

Competing Platforms and Transport Equilibrium

Nicola Rosaia¹

Columbia Business School

June 2024

ABSTRACT. I study whether platform competition in ride-hailing generates waste and whether efficiency can be enhanced by consolidating competing networks, considering the tension between market power and network economies. I construct a comprehensive dataset documenting the operations of two large platforms in New York City and use it to discipline a spatial equilibrium model of ride-hailing markets where platforms set prices strategically. Comparing the status quo with counterfactuals, I find that market power and missed network economies generate a waste of 9% and 15% of driver hours, respectively. Consolidation achieved through a merger would improve efficiency but adversely affect riders through higher prices, especially in high-density areas. In contrast, removing barriers to simultaneous multi-homing would improve efficiency and lead to lower prices, higher surplus for riders, lower traffic, and higher profits for platforms.

¹This paper is a revised version of the first chapter of my PhD dissertation at Harvard University. I am grateful to my advisors Myrto Kalouptsi, Robin Lee and Ariel Pakes. Thanks also to Giulia Brancaccio, Allan Collard-Wexler, Alessandro Gavazza, Gautam Gowrisankaran, Alessandro Lizzeri, Allen T. Zhang, three anonymous referees and numerous seminar participants for helpful comments.

1. INTRODUCTION

In recent years, online platforms have fundamentally reshaped numerous industries, including retail, hotels, media, advertising, and the taxi industry. In these markets, competition between platforms often comes at the expense of network economies, or economies of scale. Users benefit from the expansion of platforms' networks, gaining access to a larger pool of potential trading partners. If there are costs associated with users joining multiple networks simultaneously, platform competition can lead to the fragmentation of the user base across separate networks, which is inefficient. To what extent does this lead to waste? Is competition beneficial, or could efficiency be enhanced by consolidating all users into a single network? I explore this question in the ride-hailing industry focusing on New York City, the largest ride-hailing market in the United States.

I construct a comprehensive dataset of prices, wait times, drivers' compensations and movements, and match this dataset with publicly accessible trip records, providing a complete picture of the operations of two large platforms over a two-month period. Notably, the dataset tracks the movements of all idle drivers across a vast geographical area, accounting for 75% of all NYC trips. Through this data, I document that drivers spend almost half of their time empty, either idling between consecutive trips or en-route to pick-ups. The average driver utilization, defined as the share of time drivers spend transporting passengers relative to their total working hours, is only 53%. Ride-hailing vehicles account for a significant portion of traffic in Manhattan's most congested regions, with empty vehicles contributing substantially.

On one hand, a lack of interoperability contributes to this waste: passengers can only be matched with drivers who are available on the same network at the exact moment they make a request; a passenger requesting a trip on one platform cannot be matched with a driver who is online on another platform, and vice versa. By analyzing drivers' trajectories, I document that drivers are simultaneously available on both platforms for less than 5% of their idle time. This means that, even though they may frequently switch apps, drivers are mostly available on one a single platform at any given time. This leads to inefficient matching, as passengers often miss potential matches with nearby drivers who are active on a different platform, and vice versa, leading to wasted driver miles.

On the other hand, not all empty miles are necessarily wasteful. To some extent, low utilization is inherent to the nature of the taxi industry and of transportation markets in general.¹ Also, consolidating competing networks does not necessarily reduce waste, as it might lead platforms to increase prices, potentially offsetting any efficiency gains.

To separately quantify these forces, I develop a spatial equilibrium model of ride-hailing markets where platforms set prices strategically, and simulate equilibria under alternative market structures. The model focuses on the equilibrium play between two platforms that set origin-destination specific fees charged to passengers and compensations paid to drivers, while capturing two central features of ride-hailing markets. First, the market is two-sided: users respond not only to prices but also to the number of users joining on the opposing side. This is because drivers benefit from receiving more trip requests, while a higher density of idle drivers ensures that trip requests can be quickly dispatched by nearby drivers, reducing passengers' wait times before pick-up. Importantly, this spatial matching process becomes more efficient as more users participate, leading to economies of scale. Second, the model captures the spatial and dynamic elements that lead to empty driver miles. Trips relocate drivers from origin to destination, and travel patterns are asymmetric, hence drivers spend time re-positioning between consecutive trips. Although users multi-home - passengers compare price quotes across platforms before making a request, and drivers frequently switch platforms based on market conditions - drivers are only available on a single platform at any given time.

With the exception of multi-homing, the description of passengers' and drivers' behavior and their interactions is akin to widely-used spatial equilibrium models of transportation, such as Buchholz (2022), Brancaccio et Al. (2020; 2023), Castillo (2023). The framework extends this literature by incorporating platforms' endogenous price-setting behavior. Platforms use prices as instruments to coordinate the

¹Before the rise of app-based services, the average utilization of medallion taxi cabs in NYC was 42% (Buchholz, 2022). To compare this with current figures, consider that medallion taxis primarily operated in Manhattan's densest regions, where utilization rates for app-based drivers now range between 56% and 60%. Numbers are comparable in oceanic transportation, where it has been documented that dry bulk carriers travel empty 42% of the time (Brancaccio et al., 2020).

spatial search of riders and drivers, internalizing how users' behavior and the interaction between opposing sides affect market outcomes. The equilibrium concept extends previous models of two-sided platforms, such as Rysman (2004) and Armstrong (2006), to accommodate dynamics and complex forms of multi-homing.

The solution approach centers around two elements. First, I derive a Spence (1975) representation for platforms' pricing problem, which shows that platforms behave as if they were selecting a spatial allocation of drivers and passengers, while internalizing how they must set prices to induce users to behave accordingly. This allows to reframe platforms' problem focusing on the direct selection of allocations rather than prices. Second, I derive an intuitive and easy-to-compute expression for the gradient of platforms' objective function. This enables to compute the model equilibria through a gradient-based search for the allocations that maximize platforms' profits conditioned on the rival's strategy, and provides a clear characterization of the profit-maximizing prices, revisiting classic insights from two-sided markets in a spatial context.

The model is built on four key primitives, estimated by combining variation from different sources. The first primitive is the relationship between the density of idle drivers and passengers' wait times, which is estimated using granular data on driver density, wait times, and traffic speeds. The second is passengers' value of time, that is, how much they are willing to pay to reduce the time waited. This is estimated from high-frequency data on how passengers switch between platforms in response to changes in relative prices and wait times.

On the drivers' side, labor supply adjusts on both the intensive and extensive margins. For the intensive margin, I estimate a dynamic model of driver entry, exit, and movement decisions, leveraging variation in drivers' wages and labor supply to infer their outside options. Drivers are forward-looking and make optimal dynamic decisions under correct expectations on the evolution of prices and their matching probabilities. Solving this problem determines drivers' weekly working hours and the frequency with which they work across different platforms and visit various locations. On the extensive margin, there is free entry of drivers into the NYC labor pool. Intuitively, if drivers' net weekly revenues when making optimal dynamic decisions are high, more drivers join, which increases competition and lowers weekly revenues, and

vice versa. This mechanism allows the number of drivers to adjust in equilibrium, while weekly revenues net of outside options remain constant.²

The last key primitive is the extent to which platforms are willing to exercise their market power to increase profits at the expense of overall efficiency. I estimate the weight that platforms place on profits relative to total welfare by matching the observed price mark-ups with those predicted by the model.

I compare the status quo with four simulated counterfactuals. I start by computing the social planner solution in two scenarios: one where the two platforms are interoperable and one where they are not. This allows to separately quantify the extent of waste originating from two distinct sources. The first source is platforms' conduct: all else being equal, market power results in prices being set above marginal costs, leading to sub-optimal trip demand. The second is the missed network economies resulting from the lack of interoperability between platforms. I find that a social planner could increase welfare by \$173 million annually without increasing the time drivers spend on the street. Of this loss, \$53 million per year is due to market power, while the remainder is due to missed network economies. Alternatively, a social planner could reduce traffic by 24% without reducing the welfare of market participants. Platforms' conduct and missed network economies account for a waste of 9% and 15% of driver hours, respectively.

Next, I explore the extent to which these inefficiencies can be mitigated under alternative market structures. I begin by simulating the impact of a merger. Consolidation achieved through the merger improves efficiency, reducing traffic by about 10% without reducing overall welfare. However, eliminating competition leads the resulting monopolist to raise prices, exacerbating market power distortions. This negatively impacts riders, leading to an \$82 million annual drop in their surplus. The effects of the merger vary significantly by location, with low-density regions experiencing the highest efficiency gains, while high-density areas see higher price increases and greater decreases in ridership and rider surplus.

²In particular, this implies that drivers' surplus is held fixed in counterfactuals. The Supplemental Material provides a robustness analysis relaxing this assumption, which yields similar results. This robustness is due to the fact that the estimated drivers' surplus at the status quo is small. Therefore, even though it might change in percentage terms, the absolute changes are minor and overshadowed by changes in riders' surplus.

Finally, I explore whether a market configuration exists that can achieve the benefits of interoperability without sacrificing the benefits of competition. In practice, consolidating users in a single network does not necessarily require a platform merger. It can alternatively be accomplished by eliminating obstacles that prevent drivers from working simultaneously on different apps, such as technological barriers preventing third-party apps from interconnecting with multiple platforms. In the last counterfactual, I simulate a scenario where two interoperable platforms compete: each platform sets prices independently, yet idle drivers are always available to dispatch trip requests from both platforms. The outcomes are comparable to those of a merger: traffic decreases by 7.5% and overall welfare rises by \$100 million per year. The main difference is that a significant portion of these efficiency gains are passed on to riders through reduced prices. Interestingly, platforms' profits also increase, though by a smaller margin compared to the merger scenario.

Related Literature. This paper contributes to the rapidly growing empirical literature on transportation with new methods, new data, and a different scope.³ First, while previous studies treat prices (or pricing rules) as fixed model primitives, this paper introduces a spatial equilibrium model where prices are set endogenously by profit-maximizing firms, along with new methods to solve it. Second, rather than examining a single transport network in isolation, it constructs a new dataset that allows for the study of interactions between multiple networks. Third, while prior research focused on the incentives of transport network users (e.g. passengers and drivers) and their responses to exogenous policy changes (e.g. pricing), this paper emphasizes the strategic decisions and interactions of the platforms themselves.

³Within the context of urban transportation, a series of recent papers studies efficiency issues. Frechette et al. (2019) and Buchholz (2022) both study search frictions and regulation in taxi markets; Shapiro (2018) and Liu, Wan, and Yang (2021) study welfare improvements from centralization; several papers study pricing issues (e.g. Castillo (2023), Buchholz et al. (2024), Bian (2020), Ma et al. (2022), Besbes et al. (2021)) and asymmetric information (Gaineddenova (2022)). In oceanic transportation, Brancaccio et al. (2020) explore the role of the transportation sector in world trade, while Brancaccio et al. (2023) study efficiency and optimal policy in the presence of search frictions. See also Ghili and Kumar (2021) on demand and supply imbalances in ridesharing platforms; Ostrovsky and Schwarz (2018) on carpooling and self-driving cars; Kreindler (2023) on congestion pricing; Cao et al. (2021) on competition in bike-sharing; Kreindler et al. (2023) and Almagro et al. (2024) on public transportation systems; Yang (2022) on trucking markets.

As a result, this paper also contributes to the extensive empirical literature on platform economics and network effects.⁴ The importance of network economies has been well recognized in several industries, including radio and television (Besen and Johnson, 1986, Farrell et al., 1992), operating systems (Bresnahan, 2002), credit cards (Rochet and Tirole, 2003), and others. Quintessential examples are telecommunications (Brock, 1981 and Gabel, 1991) and the internet (Lehr, 1995), where interoperability standards have emerged to facilitate interconnections between separate networks. This paper argues that such standards are currently absent in ride-hailing and identifies the main channels through which this leads to waste, quantifying the resulting economic losses and the trade-offs of alternative interventions.

2. MARKET AND DATA

The US app-based taxi market, often referred to as ride-hailing, is controlled by Uber and Lyft. In 2019, these companies had about one million active drivers in the US, and about 36% of surveyed adults said they had ever used their service.⁵ I focus on New York City, one of the largest markets worldwide. The analysis excludes medallion (yellow) taxi cabs, which operate under a distinct institutional environment.

Since Uber introduced its flagship service UberX in 2012, subsequently joined by Via in 2013, Lyft in 2014, and Juno in 2016, the NYC market has experienced rapid expansion and reorganization. App-based services have quickly taken over the traditional yellow taxi cab business, accounting for about 76% of all for-hire vehicle trips as of 2019. Juno and Via ceased their operations in 2019 and 2021, leaving Uber as the dominant platform, and Lyft as its sole competitor. As of 2019, these two platforms dispatched more than 15 million trips and about 80000 drivers per month. If drivers were classified as employees, Uber alone would be the largest for-profit private employer in NYC (Parrott and Reich (2018)). Market shares have been remained more or less constant since 2019, at around 75% and 25%.

⁴E.g. Katz and Shapiro (1985), Rysman (2004), Clements and Ohashi (2005), Kaiser and Wright (2006), Argentesi and Filistrucchi (2007), Chandra and Collard-Wexler (2009), Corts and Lederman (2009), Genakos and Valletti (2011), Lee (2013), Jeziorski (2014). References include handbook chapters by Farrell and Klemperer (2007) and Jullien, Pavan, and Rysman (2021), and a survey article by Evans and Schmalensee (2005).

⁵www.pewresearch.org/fact-tank/2019/01/04/more-americans-are-using-ride-hailing-apps/

2.1. Data construction. I obtain high-frequency granular data for two large platforms, referred to as Platform 1 and Platform 2. The data spans from mid-May 2019 to mid-July 2019 and covers the vast geographical area represented in Figure 2.1. The selected area includes all of Manhattan and sizable portions of Brooklyn, Queens, and the Bronx, accounting for about 75% of all NYC ride-hailing trips.

First, I use the trip records published monthly by the NYC Taxi and Limousine Commission (TLC). This dataset contains a record of every trip originating or terminating within NYC, including information on the origin and destination, time of pick-up and drop-off, trip time and distance. I complement this with a novel comprehensive dataset of passenger prices, wait times, drivers’ compensations and movements, collected through the platforms’ Application Programming Interfaces (APIs).

APIs encode the structure of communications between a platform’s server and a user’s smartphone. Every time a user opens the app, their phone sends a message to the server with information including their desired origin and destination. The server responds with a JSON-encoded list of information that includes the upfront price the user must pay for the trip, their expected wait time before a driver can pick them up, a list of the nearest eight drivers available (represented by unique identifiers and GPS coordinates with timestamps tracing their recent trajectories), estimates of the trip’s time and distance, and a surge multiplier. This surge multiplier can be used to compute drivers’ expected compensation for the trip.⁶

I collect this data throughout the dense grid of observation points represented in Figure 2.1. Each point makes a data request every few seconds during the entire sample period, with the destination rotating among a fixed grid of locations, also represented in Figure 2.1. The density of the grid ensures that every driver is always recorded by at least one observation point, resulting in a dataset that records all drivers’ trajectories within the sample region.⁷ Importantly, drivers are only recorded when they are idle and available to receive trip requests, disappearing from the dataset when they either stop working or get matched with a passenger.

⁶In 2019, drivers were compensated according to the formula:

$$\text{driver compensation} = \text{surge multiplier} \cdot (\alpha \cdot \text{time} + \beta \cdot \text{distance})$$

where the coefficients α and β can be observed from the driver apps.

⁷To ensure that all drivers are recorded, the space between two points is kept less than the minimum distance to the furthest car observed over a two-week experimentation period.

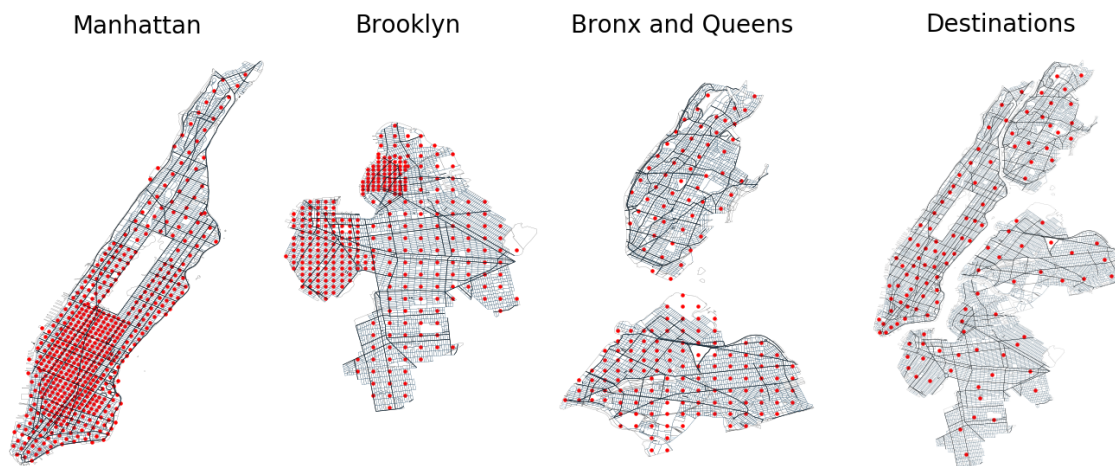


Figure 2.1. The first three panels show the observation points in different boroughs. Each point makes a data request every few seconds, with the destination rotating among the grid of destinations in the fourth panel.

I use the observed trajectories to compute the total time drivers spend idle during every one-hour window throughout the sample period. I also compute the total time drivers spend transporting passengers by summing the durations of all trips, and estimate the time drivers spend en-route to pick up passengers by summing all trips' expected wait times.⁸ These calculations provide estimates of the average number of drivers performing each task during every one-hour window in the sample.⁹ Summing across all activities yields an estimate of the total time drivers spend working, or an estimate of the average number of active drivers.

2.2. Summary statistics. The summary statistics are presented in Table 1. Platforms dispatch 81 trips per hour per square kilometer on average. Passengers typically pay \$19 per trip and wait about 3.5 minutes before pickup, while drivers receive an average of \$12. The price difference is distributed among taxes, tolls, and platforms' commissions. There are on average 42 drivers per square kilometer on the street, with 15 idle, 5 en-route to pick up a passenger, and 22 actively transporting a passenger to their destination. Driver utilization is defined as the share of time drivers spend

⁸Whenever a passenger is waiting to be picked up, a driver is traveling to pick them up

⁹For example, if drivers spend two hours idle during a one-hour interval, this indicates that, on average, there are two idle drivers

transporting passengers relative to their total working hours. This stands at 56% for Platform 1 and 46% for Platform 2. Notably, the two platforms offer nearly identical prices and pay drivers similarly, even though Platform 1 has a substantially larger market share, accounting for 75% of trips and 70% of drivers. Platforms’ wait times are also similar. In contrast, there is considerable variation in prices, wait times, driver utilization, and market activity, both spatially and over time.

	Price (\$)	Driver price (\$)	Wait time (min.)	Trips per hour (per km ²)	Active drivers (per km ²)	Idle drivers (per km ²)	Drivers en-route (per km ²)	Utili- zation
Platform	1 - 2	1 - 2	1 - 2	1 - 2	1 - 2	1 - 2	1 - 2	1 - 2
Average	19.4 - 19.5	12.5 - 12.6	3.4 - 3.7	59.9 - 21.0	29.6 - 12.7	9.6 - 5.6	3.4 - 1.3	.56 - .46
1am - 5am	18.4 - 19.0	11.9 - 12.1	3.2 - 3.5	28.1 - 11.5	13.0 - 6.9	5.1 - 3.5	1.5 - 0.7	.49 - .39
5am - 9am	20.2 - 20.5	12.9 - 13.2	3.9 - 4.4	40.3 - 12.5	20.2 - 8.3	6.6 - 3.9	2.6 - 0.9	.54 - .42
9am - 1pm	19.4 - 19.0	12.3 - 12.3	3.2 - 3.6	59.3 - 20.8	32.4 - 14.0	12.1 - 6.8	3.2 - 1.3	.53 - .43
1pm - 5pm	20.0 - 19.7	12.8 - 12.8	3.5 - 3.9	64.7 - 22.1	34.5 - 14.3	10.8 - 6.1	3.8 - 1.4	.58 - .47
5pm - 9pm	19.2 - 19.2	12.5 - 12.4	3.3 - 3.5	87.1 - 29.1	42.6 - 17.2	12.8 - 7.3	4.9 - 1.7	.58 - .48
9pm - 1am	19.2 - 19.7	12.2 - 12.6	3.3 - 3.4	79.9 - 30.1	34.9 - 15.2	10.3 - 5.8	4.3 - 1.7	.58 - .51
Lower Man.	20.8 - 21.3	13.1 - 13.1	3.4 - 3.7	174.4 - 59.8	81.3 - 33.5	21.2 - 12.7	10.0 - 3.7	.62 - .51
Midtown	20.4 - 21.2	12.6 - 12.8	3.5 - 3.7	219.1 - 68.1	110.5 - 43.6	33.5 - 19.5	13.0 - 4.2	.58 - .46
Upper Man.	20.6 - 20.4	12.7 - 12.7	3.1 - 3.5	95.3 - 30.4	46.2 - 17.6	14.3 - 7.2	4.9 - 1.8	.58 - .49
Harlem	18.2 - 18.5	12.3 - 12.5	3.2 - 3.8	53.0 - 17.1	24.0 - 9.5	7.4 - 3.9	2.8 - 1.1	.57 - .47
Bronx	16.0 - 17.4	10.9 - 12.0	3.5 - 4.3	37.6 - 6.1	17.2 - 4.5	6.2 - 2.5	2.2 - 0.4	.51 - .34
Brooklyn	19.1 - 18.1	12.7 - 12.3	3.4 - 3.6	39.0 - 20.2	19.8 - 11.0	6.6 - 4.3	2.2 - 1.2	.55 - .50
Queens	18.2 - 17.9	11.9 - 11.9	3.1 - 3.4	22.2 - 7.4	12.9 - 6.1	6.0 - 3.7	1.2 - 0.4	.45 - .32

Table 1. Summary statistics

2.3. Multi-homing. Ride-hailing users multi-home: passengers often compare price quotes across different apps before making a request, and drivers frequently switch between apps. However, platforms do not communicate with each other: a passenger requesting a trip on one platform cannot be matched with a driver who is online on

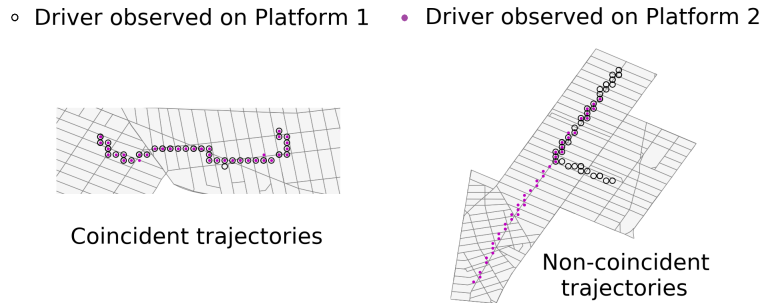


Figure 2.2. Examples of coincident and non-coincident trajectories.

another platform. In networking jargon, the platforms’ networks are not interoperable. The lack of interoperability may lead to inefficient matching, as passengers might miss potential matches with nearby drivers who are active on a different platform, and vice versa. This holds regardless of how frequently users switch platforms, since passengers can only be matched with drivers who are available on the same network at the exact moment they make a request. This inefficiency is mitigated only if most drivers run both platforms’ apps simultaneously. I refer to this type of behavior as simultaneous multi-homing.

While there are accounts of sophisticated multi-homing behavior among drivers, this tends to be the exception rather than the rule. According to TLC data, only 50% of drivers dispatched a trip for both platforms in July 2019, with the remaining half not multi-homing at all. Allon et al. (2023) further observe that drivers switch platforms between consecutive trips only about 5% of the times, which suggests 13% for the share of drivers running both apps simultaneously.¹⁰ This is likely an overestimate, since operating multiple apps concurrently requires more effort than just switching between trips. Using drivers’ geolocations, I compute a more accurate estimate.

The final data set contains records of about 37 million driver trajectories, where the average trajectory lasts about 10 minutes, and tracks the driver’s exact geolocation every 5 seconds. Intuitively, when a driver runs both apps simultaneously, the data records two trajectories on different platforms that are coincident in time and space. To account for potential data inaccuracies, I aggregate time into 20-second windows

¹⁰The two platforms dispatch about 75% and 25% of all trips, respectively. Hence, if all drivers ran both platforms simultaneously, they would switch $.75 \cdot .25 + .25 \cdot .75 = 37.5\%$ of the times. This yields an estimate of $.05/.375 = 13\%$ for the share of drivers simultaneously multi-homing.

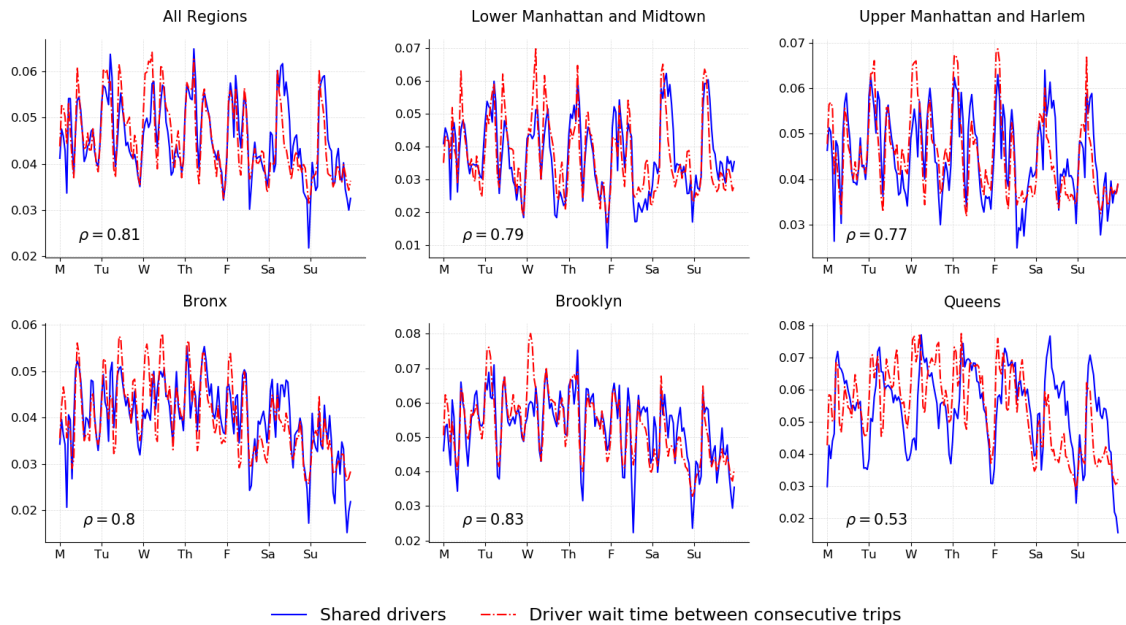


Figure 2.3. The solid line represents the share of idle drivers who work for both platforms simultaneously. The dashed line indicates the average time drivers must wait between consecutive trips, shifted and re-scaled so that the mean and variance of the two series are equal.

and space into squares of 200 squared meters, and represent each trajectory by means of the sequence of the squares visited during different time windows. Two trajectories are deemed coincident if their intersection covers at least 90% of the shortest one. Figure 2.2 shows two instances of trajectories deemed coincident and non-coincident. Overall, this test matches 4.5% of drivers.¹¹ Figure 2.3 shows how this number varies. To summarize, simultaneous multi-homing among drivers is negligible. There are several explanations for this fact. First, Figure 2.3 shows that fluctuations in the number of shared drivers are largely explained by variation in drivers' wait between consecutive trips. This is intuitive, as the main advantage for drivers to run both apps is to reduce waiting. This could partly explain why this behavior is limited

¹¹Given the density of drivers in certain regions, this test might deem as coincident trajectories of different drivers. To verify the extent of this bias, I also run a placebo test matching trajectories observed on the same platform. Given that each platform assigns its drivers a unique identifier, different trajectories on the same platform always belong to different drivers. This placebo test matches about 1.7% of drivers.

in NYC, where the average wait between trips is only about 10 minutes.¹² Second, drivers face technical challenges and incentives that discourage multi-homing. Apps often prevent drivers from receiving trip requests unless they are open on the main screen, which means that simultaneous multi-homing can only be performed using two separate phones, or with the aid of third-party software. However, technological obstacles often prevent third-party apps from functioning correctly, and subscriptions are costly.¹³ Moreover, platforms set incentives, such as loyalty payments and ratings, that increase the opportunity costs of multi-homing.¹⁴

2.4. Empty seats, full streets?¹⁵ NYC drivers spend almost half of their time empty, either idling while waiting for requests or en-route to pick up passengers. To put this into perspective, there are on average about 155 drivers per square kilometer in Midtown, 70 of whom are unoccupied. This means 3.5 drivers and 1.5 empty drivers per block, and these figures escalate during rush hours.

The level of utilization is ultimately an outcome of platforms' price-setting behavior: all else being equal, when prices are higher demand is lower, and hence utilization is lower. If platforms set prices higher than socially optimal, this results in lower than optimal utilization. As discussed in the previous section, the lack of interoperability may also lead to matching inefficiencies, wasting driver miles. However, it is important to note that low utilization does not necessarily indicate inefficiency. To some extent, empty miles are inherent to the nature of the ride-hailing industry.

First, demand is highly volatile, while drivers work longer shifts, often sitting idle between peaks in demand. Second, travel patterns are asymmetric. During an average

¹²This was confirmed by many drivers I interviewed on this subject. Most drivers mentioned minimal wait times and satisfaction with current earnings as the main reasons for working with only one platform at a time. Several drivers also cited technical difficulties and incentives described below.

¹³The largest third-party app allowing drivers to simultaneously multi-home, Mystro, charged \$5 per week for this service in 2019. According to Mystro's CEO, platforms "do not provide any official ways to interoperate with their network outside their first-party app [...] and they actively put barriers to prevent other apps from accessing their network" (authorized quote from a private conversation).

¹⁴For instance, drivers must accept at least 85% of trip requests to maintain a "pro" status which brings important benefits, and they receive bonuses for dispatching consecutive trips.

¹⁵The title of this section references a report in which a former NYC Department of Transportation's first Deputy Commissioner for Traffic and Planning advocates "for the City or State to mandate that Uber, Lyft and other TNCs limit the time that their drivers spend waiting for their next trip request" as a way to reduce Manhattan traffic congestion (Schaller, 2017).

weekday morning rush, Midtown experiences a net inflow of about 1700 trips per hour, meaning that there are 1700 more passengers traveling to Midtown than leaving it. As a result, 1700 drivers per hour cannot receive new requests from Midtown after drop-offs and must spend time idle relocating to other regions. This situation reverses during the evening rush. This pattern holds true across most regions in NYC: at any given time, either the inflow of trips substantially exceeds the outflow, or vice versa, leading to idling time and low utilization. Third, a higher number of idle drivers ensures that passengers are matched with nearby drivers.¹⁶ This improves service quality by reducing wait times for passengers and minimizes the time drivers spend en-route to pick-ups, thus reducing empty miles. In contrast, operating at excessively high utilization is inefficient and can lead to disruptions in extreme cases.¹⁷

To summarize, assessing whether low utilization is wasteful requires accounting for the market features leading to empty driver miles. These considerations inform the model presented in the next section.

3. MODEL

This section presents a model focusing on the equilibrium play between two platforms (firms, or marketplaces), $m = 1, 2$. Platforms compete by setting prices

$$p_{xij}^m, r_{xij}^m \forall i, j, x.$$

That is, they choose how much passengers pay (p_{xij}^m) and how much drivers are paid (r_{xij}^m) for trips from each origin i to each destination j . Origins and destinations $i, j \in \{1, 2, \dots, l\}$ are the regions depicted in Figure 5.2 in Section 5.2, and prices are adjusted depending on the weekday-hour $x \in \{(\text{Monday}, 0\text{am}) \dots (\text{Sunday}, 11\text{pm})\}$. I refer to x as an hour, and denote by x^{-1} and x^{+1} the hours before and after x .¹⁸

Time is continuous and measured in hours. For tractability, the transition from x to x^{+1} occurs stochastically at Poisson rate 1 - i.e. once every hour on average. I

¹⁶Section 4.3 shows that, after controlling for traffic speed and region-specific fixed effects, doubling the density of idle drivers cuts wait times in half.

¹⁷Castillo et al. (2023) note that high utilization sometimes leads to negative feedback cycles known as Wild Goose Chases: drivers travel long distances to pick up passengers, which decreases the density of available drivers, thereby increasing pick-up distances even further.

¹⁸E.g. if $x = (\text{Sunday}, 11\text{pm})$ then $x^{-1} = (\text{Sunday}, 10\text{pm})$ and $x^{+1} = (\text{Monday}, 0\text{am})$.

focus on stationary equilibria where all relevant variables remain constant within an hour x , following a cyclical pattern that repeats identically every week. Therefore, variables are indexed by x , and time is otherwise excluded from the notation.¹⁹

This is a simpler problem than the one platforms solve in real-life, mainly because origins and destinations are coarse, and prices only change on an hourly basis. Theoretically, it is straightforward to allow platforms to condition prices on additional information by expanding x to include more variables. In practice, conditioning prices on all the relevant information is computationally intractable, and platforms use heuristics that reduce dimensionality. The model simplifies by abstracting from short-lived fluctuations and instead focuses on four main ingredients.

First, the market is two-sided: riders and drivers respond not only to prices but also to the number of users on the other side. Second, the model accounts for the factors causing a temporal and spatial mismatch between riders and drivers discussed in Section 2.4. Third, the spatial matching process between riders and drivers becomes more efficient as more users join the market, resulting in economies of scale. Fourth, platforms set prices strategically, internalizing how users' behavior and the interaction between opposing sides affect market outcomes. I begin with a detailed overview of these main elements.

3.1. Main ingredients.

3.1.1. *Riders and drivers.* Passengers request trips from i to j on Platform m at Poisson rate q_{xij}^m , measured in trip requests per hour. Demand is strictly decreasing in own price p_{xij}^m and wait time w_{xi}^m , and strictly increasing in rival's price p_{xij}^{-m} and wait time w_{xi}^{-m} :

$$q_{xij}^m = \mathbf{q}_{xij}^m(p_{xij}^m, p_{xij}^{-m}, w_{xi}^m, w_{xi}^{-m}). \quad (3.1)$$

Wait times measure the amount of time a passenger must wait to be picked up, which depends on both demand and drivers' labor supply, as detailed below. To maintain consistency in time units across different variables, wait times are measured as a fraction of one hour, even though passengers typically wait for only a few minutes.

¹⁹The stationarity constraints are introduced formally in Section 3.2.2 below.

Equivalently, passenger prices are pinned down by passengers' inverse demand curves, obtained by inverting the demand curves holding wait times fixed:

$$p_{xij}^m = \mathbf{p}_{xij}^m(q_{xij}^m, q_{xij}^{-m}, w_{xi}^m, w_{xi}^{-m}). \quad (3.2)$$

Similarly, when the rival's price is held fixed, inverting $q_{xij}^m(\cdot)$ in its first argument yields m 's residual inverse demand curve. Slightly abusing notation, I denote this by

$$p_{xij}^m = \mathbf{p}_{xij}^m(q_{xij}^m, p_{xij}^{-m}, w_{xi}^m, w_{xi}^{-m}). \quad (3.3)$$

This allows to define the weekly gross surplus of m 's passengers W^m , measured in dollars per week, as the total area below these curves across all origins, destinations, and weekday-hours. The net surplus CS^m is the gross surplus minus the prices paid:

$$W^m \equiv \sum_{xij} \int_0^{q_{xij}^m} \mathbf{p}_{xij}^m(z, p_{xij}^{-m}, w_{xi}^m, w_{xi}^{-m}) dz; \quad CS^m \equiv W^m - \sum_{xij} q_{xij}^m p_{xij}^m. \quad (3.4)$$

On drivers' side, labor supply adjusts on two distinct margins. First, it adjusts on the extensive margin, which determines the number of potential workers, denoted by N . This can be thought of as the pool of qualified individuals who register to drive in NYC, which entails clearing a set of administrative fees and procedures. I will refer to these potential workers simply as drivers, with the understanding that, at any given time, a driver may not necessarily be working. Indeed, as detailed in Section 3.2 below, drivers solve a dynamic problem, frequently choosing whether to work or remain inactive, which platform to work for, and where to position in the city network. They respond to the system of driver prices r , wait times w , and the frequencies of trip requests they receive while idle.²⁰ I refer to these frequencies as drivers' matching rates, which are given by the ratios between demand and the number of competing available drivers:

$$\theta_{xij}^m = \boldsymbol{\theta}_{xij}^m(q_{xij}^m, s_{xi}^m) \equiv \frac{q_{xij}^m}{s_{xi}^m}. \quad (3.5)$$

Drivers' dynamic decisions determine n_x , the average number of drivers working in hour x , as well as s_{xi}^m , the average number of drivers who are available on Platform

²⁰I omit indices when denoting full vectors, e.g. $r = [r_{xij}^m \forall x, ij m]$.

Again, to ensure consistency in time units, matching rates are measured in matches per hour, even though drivers typically wait only a few minutes before receiving a request.

m - i.e. idle, waiting to be matched - in region i . These must satisfy

$$\sum_{im} s_{xi}^m + \sum_{ijm} q_{xij}^m (w_{xi}^m + t_{xij}) = \underset{\substack{\text{active} \\ \text{(labor supply)}}}{n_x} \leq \underset{\substack{\text{registered} \\ \text{(labor pool)}}}{N} \quad (3.6)$$

where t_{xij} denotes is the average duration of a trip from i to j in hour x , measured as a fraction of one hour.

Equation (3.6) is an accounting identity stating that, at any given time, each active driver is either already matched or idle and available to be matched. Intuitively, for a driver, a trip entails spending w_{xi}^m minutes traveling to pick up the passenger, and t_{xij} minutes to transport the passenger from origin to destination. Hence, $\sum_{ijm} q_{xij}^m (w_{xi}^m + t_{xij})$ is the average number of drivers who are matched, obtained by multiplying the number of trips by the duration of each trip. If there are n_x active drivers in total, this implies that the number of drivers who are idle, $\sum_{im} s_{xi}^m$, must satisfy the above equation.

Both n_x and s_{xi}^m are continuous-time quantities, measuring the number of drivers who are active and idle on average at any given time instant during hour x . Section 3.2 describes how these quantities are outcomes of a process where drivers continuously transition between being idle and being matched, enter and exit, and move between regions. Importantly, this dynamic process accounts for the factors leading to empty driver miles discussed in Section 2.4. First, although demand is concentrated around peak hours, drivers may choose to work longer shifts, creating a temporal mismatch between supply and demand. Second, travel patterns are asymmetric, and trips relocate drivers from origin to destination, causing drivers to re-position between consecutive trips. Third, drivers can multi-home by frequently switching platforms. However, regardless of how often they switch, they can only be available on one platform at a time.

Four primitives dictate how drivers react to market conditions. First, drivers incur a per-mile cost c_d for dispatching a trip. Second, while they are active - either idle or matched - they incur a flow cost c_x . This encompasses expenses like fuel, vehicle leasing or depreciation, as well as opportunity costs such as the value of leisure time. I refer to these costs as drivers' (average) dynamic outside options, which can vary over time depending on weekday-hours. A third parameter, ν , governs heterogeneity in drivers' dynamic outside options, which in turn determines the elasticity of their

labor supply on the intensive margin. That is, how many more (less) hours each driver will work following and increase (decrease) in wages.

When making optimal dynamic decisions, a driver’s weekly revenues net of these costs yield an optimal net weekly payoff u , which is a function of the payoff-relevant variables:

$$u = \mathbf{u}(r, \theta, w).$$

To ensure consistency in time units, u is measured in dollars per hour, even though it represents a weekly average. Drivers weigh u against a fourth parameter \bar{u} , encompassing (amortized) weekly registration costs, e.g. administrative fees and inspections, as well as wages from driving outside NYC. I model this through a free entry condition:

$$u = \bar{u}. \tag{3.7}$$

This captures drivers’ extensive margin labor supply decisions as well as their mobility across adjacent urban areas, similarly to worker free entry conditions typically found in spatial urban models (e.g. Ahlfeldt et al. 2015).²¹ Intuitively, if u is higher than \bar{u} more drivers will join the NYC labor pool, and competition among them will lower u , and vice versa. This pins down the number of drivers N registered in NYC. On the other hand, it implies that drivers’ surplus, i.e. their weekly net payoff, is pinned down by \bar{u} in equilibrium, hence it remains fixed in counterfactuals.

3.1.2. *Matching.* The second main ingredient is the matching technology, which is modeled by means of a function

$$w_{xi}^m = \mathbf{w}_{xi}(s_{xi}^m). \tag{3.8}$$

This accounts for spatial matching frictions and platforms’ first-dispatch protocol: when a rider requests a trip, she is immediately matched with the nearest available driver. Indeed, recall that s_{xi}^m measures the average number of drivers who are available at any given time instant during hour x . Also, because time is continuous, there

²¹The Supplemental Material includes an extension where Equation (3.7) is replaced by a constant elasticity supply curve. It also provides a robustness analysis for the counterfactuals under different calibrations for the extensive margin elasticity, yielding results similar to those obtained under Equation (3.7). This robustness is due to the fact that, as shown in Section 4, drivers’ surplus is very small across all specifications. Thus, even though drivers’ surplus might change in percentage terms, the absolute changes are small and overshadowed by changes in other variables.

is never more than one request arriving at the exact same instant. Hence, when a rider requests a trip, there are exactly s_{xi}^m drivers available to her. This justifies the function in Equation (3.8). Intuitively, a higher density of idle drivers implies that, on average, a rider will be matched to a driver closer to her location, resulting in lower wait times.

Note that wait times depend on both passenger demand and drivers' labor supply through Equation (3.6), even though these factors are not direct arguments of function $w_{xi}(\cdot)$. Holding drivers' labor supply constant, an increase in demand reduces the number of available drivers, thereby increasing wait times. Conversely, holding demand constant, an increase in drivers' labor supply leads to more available drivers, thus reducing wait times. Also, importantly, a proportional increase in both demand and labor supply results in more available drivers, thus reducing wait times. This mechanism generates economies of scale: the matching process becomes more efficient as more users join a platform.

3.1.3. *Conduct.* The third main ingredient is platforms' conduct, that is, how they strategically set prices. This involves two elements.

The first is the extent to which platforms are willing to exercise their market power to increase profits at the expense of overall efficiency. In practice, platforms may avoid setting profit-maximizing prices, both to avoid regulatory and media scrutiny and to incentivize adoption over the long term (see e.g. Castillo (2023), Gutierrez (2021), Sullivan (2022)). The model captures this in a reduced-form way by assuming that Platform m 's objective is to maximize a combination O^m of her weekly profits and the total surplus generated in her marketplace:

$$O^m = \lambda^m \Pi^m + (1 - \lambda^m)(\Pi^m + CS^m). \quad (3.9)$$

Note that drivers' surplus does not appear in this expression because it is pinned down by free entry (Condition 3.7). m 's weekly profits are the difference between the prices received from passengers and payments to drivers:

$$\Pi^m = \sum_{xij} q_{xij}^m (p_{xij}^m - r_{xij}^m).$$

When $\lambda^m = 1$, platforms maximize profits. When $\lambda^m = 0$, platforms maximize the total surplus, which is the sum of profits and net customer surplus. Intermediate values of λ^m capture the weight platforms place on profits, a parameter I estimate from data. Section 5.3 also considers counterfactuals with alternative values of λ^m .

Second, to allow platforms to set prices strategically, one needs to consider how they expect price deviations to impact market outcomes. There are multiple rational expectations that platforms might hold, since users respond not only to prices but also to wait times and matching rates, which in turn depend on users' participation decisions. This creates the standard coordination problem in two-sided markets: one side will not join if the opposing side does not join, and vice versa.

To address this, I assume that, when evaluating price deviations, platforms hold fixed not only the system of two-sided prices p^{-m}, r^{-m} but also the system of pick-up times w^{-m} and matching rates θ^{-m} prevailing on the rival platform. They consider a market outcome to be possible under p^m, r^m if it is consistent with passengers' demand curves, drivers' dynamic decisions and free entry, and how the interaction between opposing sides affects w^m, θ^m through their matching technology, holding $p^{-m}, r^{-m}, w^{-m}, \theta^{-m}$ fixed. In equilibrium, no price deviation is profitable when evaluated based on its best possible outcome. In other words, platforms cannot expect to benefit from any price deviation. This equilibrium notion aligns with previous work on two-sided platforms, which typically assumes away coordination failures.

Section 3.2 describes the supply-side dynamics in detail. With these in place, Section 3.3 provides a formal definition of the market equilibrium, and further discusses similarities and differences with previous literature on two-sided platforms. Addressing platforms' pricing problem within a dynamic model presents technical challenges, which are addressed in Sections 3.4-3.6. Given the technical nature of the remainder of this section, readers interested only in the empirical analysis may skip directly to Section 4.

3.2. Supply-side dynamics.

3.2.1. *Entry, exit and movement.* Drivers make decisions in continuous time responding to prices, wait times, and matching rates. They are forward-looking, forming correct expectations on the evolution of these variables at various locations. They have a

zero discount rate, that is, they maximize their average weekly payoff.²² Solving their dynamic decision problem yields drivers' optimal average weekly payoff u , measured in dollars per hour. Readers who are not familiar with models without discounting can simply think at u as drivers' opportunity cost of time. The optimality conditions below are similar to those that would obtain in a model with discounting; the only difference being that, instead of weighting future payoffs according to a discount rate, the opportunity cost of time is discounted every hour.²³

Drivers begin by making hourly entry and exit decisions. At the beginning of hour x , those who were inactive during x^{-1} can enter at a random location, receiving

$$V_x = \sum_i P_{i|x} V_{xi}. \quad (3.10)$$

$P_{i|x}$ denotes the probability of starting from region i , hence $\sum_i P_{i|x} = 1$, while V_{xi} denotes the inclusive value of the entry choice

$$V_{xi} = \mathbb{E} \max\{V_{in|x_i} + \epsilon_{in}, V_{out|x} + \epsilon_{out}\}. \quad (3.11)$$

In the above expression, $V_{in|x_i}$ and $V_{out|x}$ denote the continuation values of working and staying inactive, respectively. ϵ is a i.i.d. preference shock drawn from a type-I extreme value distribution with scale parameter γ .

The distribution of $\epsilon_{out} - \epsilon_{in}$ captures heterogeneity in drivers' outside options, hence γ governs the degree of heterogeneity, which determines the sensitivity of entry decisions to differences in continuation values. The probability that a driver chooses to enter, denoted by $\sigma_{in|x_i}$, satisfies the logit formula

$$\sigma_{in|x_i} = (\exp \gamma^{-1} V_{in|x_i} + \exp \gamma^{-1} V_{out|x})^{-1} \exp \gamma^{-1} V_{in|x_i}. \quad (3.12)$$

Similarly, I denote the probability that a driver opts to stay inactive by

$$\sigma_{out|x_i} = 1 - \sigma_{in|x_i}. \quad (3.13)$$

²²The absence of discounting is a natural assumption, since drivers make frequent decisions in an environment that repeats cyclically every week. From an empirical standpoint, this is indistinguishable from a model with a small positive discount rate.

²³A more detailed explanation can be found in the Supplemental Material (Rosaia, 2024), which contains a self-contained exposition of drivers' problem with discounting, and derives the optimality conditions below in the limit as the discount rate vanishes.

Drivers who stay inactive discount the opportunity cost of time for one hour, and face the entry choice again in $x+1$, hence they receive

$$V_{\text{out}|x} = -u + V_{x+1}. \quad (3.14)$$

Drivers who enter choose which platform to work for, receiving continuation value

$$V_{\text{in}|xi} = \mathbb{E} \max\{V_{m|xi} + \epsilon_m : m = 1, 2\} \quad (3.15)$$

where $V_{m|xi}$ denotes the continuation value of working for Platform $m = 1, 2$, and ϵ is a i.i.d. preference shock drawn from a type-I extreme value distribution. The probability that a drivers chooses Platform m , denoted by $\sigma_{m|xi}$, takes the logit formula

$$\sigma_{m|xi} = \left(\sum_m \exp V_{m|xi} \right)^{-1} \exp V_{m|xi}. \quad (3.16)$$

Drivers choosing to work for m remain idle until either they get matched, or they move to an adjacent region, or x transitions to $x+1$. I first define their continuation value in these three scenarios, and then explain how these contribute to $V_{m|xi}$.

First, idle drivers get matched with passengers routed towards j at rate θ_{xij}^m , receiving

$$r_{xij}^m - c_d d_{ij} - (w_{xi}^m + t_{xij})(c_x + u) + V_{j|xim}. \quad (3.17)$$

That is, matched drivers receive the compensation r_{xij}^m upfront and pay a distance-based cost $c_d d_{ij}$, which depends on the distance d_{ij} and the per-mile cost c_d . They then spend w_{xi}^m hours en-route to pick up the passenger and t_{xij} hours traveling from i to j , incurring the dynamic outside option c_x and the opportunity cost of time u for the total duration of $w_{xi}^m + t_{xij}$ hours. Upon reaching destination, they receive the expected continuation value

$$V_{j|xim} = (w_{xi}^m + t_{xij})V_{x+1j} + [1 - (w_{xi}^m + t_{xij})]V_{\text{in}|xj}. \quad (3.18)$$

Intuitively, since both wait times and travel times are measured in fractions of hours, drivers dispatched during hour x arrive at destination during $x+1$ with probability $w_{xi}^m + t_{xij}$. In this case, they choose whether to keep working or become inactive, receiving the inclusive value of the entry choice. With the remaining probability, $1 - (w_{xi}^m + t_{xij})$, they arrive before x transitions to $x+1$ and choose whether to switch platform, receiving the inclusive value of the platform choice.

Second, at Poisson rate λ , drivers idle in i can move towards a region j within J_i , which includes i and the regions adjacent to it. With some probability $P_{j|xi}$, drivers choosing j relocate to j and choose whether to switch platform, receiving the inclusive value of the platform choice; with the remaining probability, $1 - P_{j|xi}$, they remain idle in i but may still switch platforms.²⁴ Drivers choosing $j = i$ always remain in i , meaning $P_{i|xi} \equiv 1$. The inclusive value of this movement decision is given by

$$V_{\text{move}|xi} = \mathbb{E} \max\{P_{j|xi} V_{\text{in}|xj} + (1 - P_{j|xi}) V_{\text{in}|xi} + \epsilon_j : j \in J_i\} \quad (3.19)$$

where ϵ is drawn i.i.d. from a type-I extreme value distribution. Hence, drivers choose j with probability given by the logit formula

$$\sigma_{j|xi} = \left[\sum_j \exp P_{j|xi} (V_{\text{in}|xj} - V_{\text{in}|xi}) \right]^{-1} \exp P_{j|xi} (V_{\text{in}|xj} - V_{\text{in}|xi}). \quad (3.20)$$

Third, x transitions to x^{+1} at Poisson rate 1, at which point drivers choose whether to continue working or become inactive, receiving the inclusive value of the entry decision, $V_{x^{+1}i}$.

Hence, finally, idle drivers receive the expected continuation value

$$\begin{aligned} V_{m|xi} = & (1 + \lambda + \sum_j \theta_{xij}^m)^{-1} \{-c_x - u + \lambda V_{\text{move}|xi} + V_{x^{+1}i} \\ & + \sum_j \theta_{xij}^m [r_{xij}^m - c_d d_{ij} - (w_{xi}^m + t_{xij})(c_x + u) + V_{j|xim}]\}. \end{aligned} \quad (3.21)$$

In words, drivers remain idle for an average of $(1 + \lambda + \sum_j \theta_{xij}^m)^{-1}$ hours, during which they incur the dynamic outside option c_x and the opportunity cost of time u . Then, one of the following happens: with probability proportional to θ_{xij}^m , they are dispatched to j , receiving the continuation value in Expression 3.17; with probability proportional to λ , they make a movement decision, receiving the inclusive value $V_{\text{move}|xi}$; with probability proportional to 1, x transitions to x^{+1} , and they receive the inclusive value of the entry/exit decision, $V_{x^{+1}i}$.

²⁴For the empirical analysis I set $\lambda = 60/5$, which means that idle drivers face movement decisions once every five minutes on average. For consistency with travel times, I set $P_{j|xi} = (\lambda t_{xij})^{-1}$. This implies that, on average, it takes t_{xij} hours for a driver who always chooses j to relocate there.

3.2.2. *Stationary distribution.* When the number of drivers N is given, their choice probabilities σ determine the hourly distribution of drivers available on different platforms at various locations, s_{xi}^m , and drivers' overall labor supply, n_x . These must satisfy a system of recursive equations.

First, let η_{xi} be the average number of drivers who begin hour x by facing the entry choice in region i . We must have

$$\eta_{xi}^{\text{entry choice}} = \sum_m s_{x^{-1}i}^m + \sum_{jm} s_{x^{-1}j}^m \theta_{x^{-1}ji}^m (w_{x^{-1}j}^m + t_{x^{-1}ji}) + (N - n_{x^{-1}}) P_{i|x}. \quad (3.22)$$

Indeed, when x^{-1} transitions to x , on average, $\sum_m s_{x^{-1}i}^m$ drivers are idle in i and face the entry choice immediately; $\sum_{jm} s_{x^{-1}j}^m \theta_{x^{-1}ji}^m (w_{x^{-1}j}^m + t_{x^{-1}ji})$ drivers are matched towards i and face the entry choice upon drop-off; $N - n_{x^{-1}}$ drivers are inactive and face the entry choice in i with probability $P_{i|x}$.

Drivers choose to exit (or to remain inactive) with probability $\sigma_{\text{out}|xi}$, so we must have

$$N - n_x = \sum_i \eta_{xi} \sigma_{\text{out}|xi}. \quad (3.23)$$

On the other hand, drivers choose to enter (or to continue working) with probability $\sigma_{\text{in}|xi}$, and must choose which platform to work for. Let π_{xi} be the average number of drivers facing the platform choice in region i during hour x . Recall that drivers may also switch platforms upon moving, or directly upon drop-off if they reach their destination before the transition to x^{+1} (otherwise, they must first choose whether to keep working). Hence, we must have

$$\begin{aligned} \pi_{xi}^{\text{platform choice}} &= \eta_{xi} \sigma_{\text{in}|xi} + \sum_{jm} s_{xj}^m \theta_{xji}^m [1 - (w_{xj}^m + t_{xji})] \\ &\quad + \lambda \sum_{jm} [s_{xj}^m \sigma_{i|xj} P_{i|xj} + s_{xi}^m \sigma_{j|xi} (1 - P_{j|xi})] \end{aligned} \quad (3.24)$$

Indeed, $\eta_{xi} \sigma_{\text{in}|xi}$ drivers face the platform choice in i upon entry; $s_{xi}^m \theta_{xji}^m$ drivers dispatched from j to i reach destination before the transition to x^{+1} with probability $1 - (w_{xj}^m + t_{xji})$, facing the platform choice in i directly upon drop-off; λs_{xj}^m drivers make a movement decision from j during hour x , move towards i with probability $\sigma_{i|xj}$, and reposition to i with probability $P_{i|xj}$, facing the platform choice in i ; λs_{xi}^m

drivers make a movement decision from i , move towards j with probability $\sigma_{j|xi}$, and remain in i with probability $1 - P_{j|xi}$, facing the platform choice in i .

These drivers choose Platform m with probability $\sigma_{m|xi}$, remaining idle until either they get matched (at rate $\theta_{xi}^m \equiv \sum_j \theta_{xij}^m$) or make a movement decision (at rate λ), before x transitions to x^{+1} . The average number of idle drivers is the difference between the average inflow and the average outflow.

$$s_{xi}^m = \underbrace{\pi_{xi}\sigma_{m|xi}}_{\text{inflow}} - \underbrace{(\theta_{xi}^m + \lambda)s_{xi}^m}_{\text{outflow}}. \quad (3.25)$$

Finally, Equation (3.6) must hold. Replacing q_{xij}^m with $s_{xi}^m\theta_{xij}^m$, this can be written as

$$\sum_{im} s_{xi}^m + \sum_{ijm} s_{xi}^m \theta_{xij}^m (w_{xi}^m + t_{xij}) = n_x. \quad (3.26)$$

That is, at any given time, active drivers are idle or matched, and the number of matched drivers is obtained by multiplying the average number of matches, $s_{xi}^m\theta_{xij}^m$, by the average time needed to dispatch a trip, $w_{xi}^m + t_{xij}$.

3.2.3. Optimality. Conditions (3.22)-(3.26) describe the stationary distributions that are compatible with the entry, exit and movement decisions of N drivers, encapsulated by the optimal choice probabilities σ .

Definition 1. s is consistent with the optimal entry, exit, and movement decisions of N drivers if Conditions (3.10)-(3.26) are satisfied.

Drivers' dynamic decisions also determine their average weekly payoff u . It can be shown that this admits an intuitive expression.

Proposition 1. *If Conditions (3.10)-(3.26) are satisfied then*

$$u = \left\{ \sum_{xijm} s_{xi}^m \theta_{xij}^m [r_{xij}^m - c_d d_{ij} - (t_{xij} + w_{xi}^m) c_x] - \sum_{xim} s_{xi}^m c_x + \mathcal{E} \right\} \cdot (168 \cdot N)^{-1} \quad (3.27)$$

where \mathcal{E} is a term that captures drivers' payoffs from the i.i.d. shocks.

The first term on the right-hand side is drivers' weekly total surplus, that is, the sum of weekly revenues from matches net of the costs of time and distance, minus the cost of idling time, plus a term that captures utilities from the i.i.d. shocks. Dividing this

by 168 (the number of weekday-hours) and by the number of drivers N yields the surplus u for an individual driver measured in dollars per hour.

Definition 1 describes the supply-side adjustment on the intensive margin, i.e. holding N fixed. Free entry further requires $u = \bar{u}$ (Condition 3.7). When this is also satisfied, I say that s is driver-optimal.

Definition 2. (i) s is driver-optimal if Conditions (3.7) and (3.10)-(3.26) are satisfied.
(ii) s^m is driver-optimal if there exists s^{-m} such that s is driver-optimal.

Condition (i) defines a supply curve, describing how drivers' distribution responds to changes in r , θ and w accounting for adjustments on both the intensive and extensive margins. Condition (ii) defines the supply curves faced by individual platforms.

3.2.4. *Inversion.* Every distribution of available drivers corresponds to an essentially unique system of payments drivers must receive. This forms the basis for identification, and supports both the theoretical results and the computational algorithm. This section makes this idea precise. Readers not interested in theory and computation may jump directly to Section 3.3 and skip Sections 3.4-3.6 that follow.

Note that drivers' optimality conditions (Equations 3.10-3.21) and average weekly payoff (Equation 3.27) depend on prices only through the mean payoffs

$$\bar{\delta}_{xi}^m = -c_x + \sum_j \theta_{xij}^m [r_{xij}^m - c_d d_{ij} - (w_{xi}^m + t_{xij})c_x]. \quad (3.28)$$

This means that different price systems that yield the same $\bar{\delta}$ result in identical choice probabilities σ and average weekly payoff u . Intuitively, this is because $\bar{\delta}_{xi}^m$ fully captures drivers' incentives for working on Platform m in region i : for each hour spent searching, they incur the outside option c_x , and get matched $\sum_j \theta_{xij}^m$ times on average, receiving the compensation net of the costs of time and distance. Slightly abusing notation, I write this as

$$u = u(\bar{\delta}, \theta, w)$$

and say that s is driver-optimal under $\bar{\delta}, \theta, w$ if it is driver-optimal under r, θ, w , with the understanding that $\bar{\delta}$ is a function of prices.

Drivers' distribution is invertible in δ : for every distribution s , there is a unique δ such that s is driver-optimal. Inversion nests two steps, corresponding to the intensive and extensive margin. First, when N is held fixed, it can be shown that there is a unique system of mean payoffs rationalizing s as the outcome of drivers' optimal dynamic decisions. I denote this by $\delta(s, \theta, w, N)$.

Proposition 2. *For every θ, w , every s strictly positive and every N large enough, there exists a unique $\delta \equiv \delta(s, \theta, w, N)$ such that s is consistent with the optimal entry, exit, and movement decisions of N drivers.*

This relates to standard identification results in dynamic discrete choice models, where it is well-known that choice probabilities identify differences in payoffs (see e.g. Magnac and Thesmar, 2002). Here, drivers' choice probabilities are not observed, as multiple systems σ can result in the same distribution s . Instead, s reveals the time drivers spend idle in various locations, which identifies δ .

Optimality further requires N to adjust until free entry is satisfied:

$$\mathbf{u}(\delta(s, \theta, w, N), \theta, w) = \bar{u}. \quad (3.29)$$

This suggests an inversion procedure that searches for an N satisfying the above equation, while adjusting δ as a function of N to ensure that Conditions (3.10)-(3.26) also hold. The Supplemental Material presents an algorithm performing these steps. Intuitively, consider increasing N while keeping s fixed. This means that more drivers contribute the same labor supply, hence drivers must spend more time inactive. For this to be the case, working must become less attractive, meaning that the payoffs in $\delta(s, \theta, w, N)$ must decrease as a function of N . This suggests that, first, $\mathbf{u}(\delta(s, \theta, w, N), \theta, w)$ should decrease with N . Second, as payoffs decrease, drivers spend more time inactive, receiving zero payoff. Hence $\mathbf{u}(\delta(s, \theta, w, N), \theta, w)$ can approach zero but not become negative. This reasoning leads to the following.

Proposition 3. *Function $\mathbf{u}(\delta(s, \theta, w, N), \theta, w)$ is strictly decreasing and continuous in N , and maps N onto $(0, +\infty)$.*

This implies that function $\mathbf{u}(\delta(s, \theta, w, \cdot), \theta, w)$ is invertible: for every \bar{u} strictly positive, Equation (3.29) has a unique solution for N . Hence there exists a unique δ such that both free entry and Conditions (3.10)-(3.26) are satisfied.

Proposition 4. *For every θ, w and every s strictly positive there exists a unique δ such that s is driver-optimal.*

That is, there exists a unique δ rationalizing s as an outcome of drivers' behavior, accounting for adjustments on both the intensive and extensive margins. I denote this by:

$$\delta = \delta(s, \theta, w). \quad (3.30)$$

When the rival's prices, hence δ^{-m} , are held fixed, a similar inversion defines the inverse supply curves faced by individual platforms.

Proposition 5. *For every δ^{-m}, θ, w and every s^m strictly positive there exists a unique δ^m such that s^m is driver-optimal.*

That is, holding the rival's prices fixed, there is a unique system of mean payoffs δ^m that m must provide to drivers in order to induce them to behave consistently with s^m . Slightly abusing notation, I denote this by

$$\delta^m = \delta^m(s^m, \delta^{-m}, \theta, w). \quad (3.31)$$

Equations (3.30) and (3.31) define the supply-side analogues of Equations (3.2) and (3.3). Therefore, I will refer to these as drivers' inverse supply curve and m 's residual inverse supply curve, respectively.

3.3. Equilibrium. When evaluating price deviations, platforms hold fixed both the system of two-sided prices p^{-m}, r^{-m} and the system of pick-up times w^{-m} and matching rates θ^{-m} prevailing on the rival platform. m deems q^m, s^m a possible outcome of p^m, r^m if this is consistent with how users respond, holding $p^{-m}, r^{-m}, w^{-m}, \theta^{-m}$ fixed.

Definition 3. Platform m deems q^m, s^m a possible outcome of p^m, r^m if (i) $q^m = \mathbf{q}^m(p^m, p^{-m}, w^m, w^{-m})$ and (ii) s^m is driver-optimal under δ, θ, w , where (iii) $w^m = w(s^m)$ and $\theta^m = \theta^m(q^m, s^m)$.²⁵

First, platforms correctly internalize their demand curves. Second, they correctly internalize drivers' behavior, accounting for adjustments on both the intensive and

²⁵In what follows I remove indices from an equation when it applies to all coordinates, e.g. $q^m = \mathbf{q}^m(p^m, p^{-m}, w^m, w^{-m})$ means $q_{xij}^m = \mathbf{q}_{xij}^m(p_{xij}^m, p_{xij}^{-m}, w_{xi}^m, w_{xi}^{-m})$ for all x, ij .

extensive margins. Third, they internalize how the interaction between opposing sides affects w^m, θ^m through their matching technology, holding w^{-m}, θ^{-m} fixed.

Definition 4. q, s, p, r is an equilibrium if, for all $m = 1, 2$, q^m, s^m, p^m, r^m solves

$$\max_{p^m, r^m} (\max_{q^m, s^m} O^m \text{ s.t. } q^m, s^m \text{ is a possible outcome of } p^m, r^m). \quad (3.32)$$

For any price system p^m, r^m , the inner problem in (3.32) defines its best possible outcome. That is, the allocation q^m, s^m maximizing m 's objective O^m (Equation 3.9) among all possible outcomes of p^m, r^m . Definition 4 states that in equilibrium no price deviation is profitable when evaluated based on its best possible outcome. In other words, platforms cannot expect to benefit from any deviation.

This equilibrium notion aligns with previous work on two-sided platforms, which typically assumes away coordination failures. The common approach assumes either that firms can directly select their users' utilities from joining (e.g. Armstrong, 2006) or that firms can directly select the number of their users (e.g. Rysman, 2004). These two notions are equivalent in a monopoly but differ when there is more than one firm. Intuitively, firms can expect to steal users from their rival upon deviating when they compete in users' utilities but not when they select directly the number of their users. This paper aligns more closely with the first notion: platforms take as given all factors affecting users' payoffs when joining the rival (prices, wait times, matching rates) while they can expect to steal both passengers and drivers from the rival upon deviating. On the other hand, this model differs from previous literature by accommodating complex forms of multi-homing. First, drivers are allowed to switch platforms dynamically. This implies that platforms' share drivers' labor pool N , and that drivers' weekly utility from joining the market, $u(r, \theta, w)$, is a function of both platforms' market outcomes, making previous approaches not directly applicable. Second, this equilibrium notion readily extends to the case where drivers simultaneously multi-home, a counterfactual considered in Section 5.

3.4. Spence's monopolist. Platforms behave as if they were selecting an allocation q^m, s^m while internalizing how they must set prices to induce their users to behave accordingly. To make this precise, the remainder of this section and Sections 3.5-3.6

below simplify the notation by omitting $p^{-m}, r^{-m}, \theta^{-m}, w^{-m}$. The understanding is that the rival's variables are held fixed, and the focus is on m 's optimal response.

First, note that platforms' objective O^m (Equation 3.9) can be written as the difference between a revenue term R^m and a cost term C^m :

$$O^m = [\lambda^m \sum_{xij} q_{xij}^m p_{xij}^m + (1 - \lambda^m) W^m] - [\sum_{xij} q_{xij}^m r_{xij}^m] \equiv R^m - C^m.$$

In words, platforms maximize a combination of revenues and rider surplus, minus total costs. In turn, C^m can be written in terms of drivers' mean payoffs δ^m :

$$C^m = \sum_{xij} q_{xij}^m r_{xij}^m = \sum_{xi} s_{xi}^m c_x + \sum_{xij} q_{xij}^m [c^d d_{ij} + (w_{xi}^m + t_{xij}) c_x] + \sum_{xi} s_{xi}^m \delta_{xi}^m.$$

²⁶Intuitively, platforms must compensate drivers for their cost of idling time and the time- and distance-based costs they incur when dispatching trips. Additionally, they must provide drivers with location-specific incentives, captured by δ^m .

When the rival's variables are held fixed, p^m , W^m and δ^m are pinned down as functions of q^m, s^m, w^m, θ^m by m 's residual inverse demand and supply curves (Equations 3.3, 3.4 and 3.31). In turn, $w^m = w(s^m)$ and $\theta^m = \theta^m(q^m, s^m)$ are functions of q^m, s^m . Hence, both the cost and the revenue term are ultimately functions of q^m, s^m :

$$R^m = \mathbf{R}^m(q^m, s^m), \quad C^m = \mathbf{C}^m(q^m, s^m).$$

Combining these observations shows that Problem 3.32 is equivalent to:

$$\max_{q^m, s^m} \mathbf{R}^m(q^m, s^m) - \mathbf{C}^m(q^m, s^m). \quad (3.33)$$

This is reminiscent of a Spence (1975) monopolist. Intuitively, wait times reflect the quality of service for passengers. Platforms can influence service quality by incentivizing drivers to spend more time idle. Hence they behave as if they were selecting quantity and quality of service, subject to a revenue function that captures passengers' demand curves and a cost function that captures drivers' behavior.

Problem 3.33 should be viewed as an "as if" representation of platforms' pricing problem, since platforms hold fixed the rival's prices p^{-m}, r^{-m} , rather than the rival's allocation q^{-m}, s^{-m} . Importantly, firms can expect to attract users from their rival

²⁶To see this, multiply both sides of Equation (3.28) by s_{xi}^m , use $s_{xi}^m \theta_{xij}^m = q_{xij}^m$ and rearrange terms.

upon deviating. Mathematically, this difference is captured by the fact that platforms internalize their residual inverse demand and supply curves (Equations 3.3 and 3.31), which take as given the rival's prices, rather than the full inverse demand and supply curves (Equations 3.2 and 3.30), which take as argument the rival's allocation.

3.5. Optimality conditions. This formulation simplifies in two ways. First, Problem 3.33 is unconstrained.²⁷ Second, and most importantly, it allows for the derivation of intuitive expressions for platforms' optimality conditions.

Proposition 6. *Take any q^m, s^m strictly positive. Let V be the system of drivers' value functions when $w^m = \mathbf{w}(s^m)$, $\theta^m = \boldsymbol{\theta}^m(q^m, s^m)$ and $\delta^m = \boldsymbol{\delta}^m(s^m, \theta^m, w^m)$, and r^m be any price system satisfying Equation (3.28). Then:*

$$\frac{dC^m(q^m, s^m)}{dq_{xij}^m} = c^d d_{ij} + (w_{xi}^m + t_{xij})(c_x + \bar{u}) + V_{m|x_i} - V_{j|x_{im}} \quad (3.34)$$

$$\begin{aligned} \frac{dC^m(q^m, s^m)}{ds_{xi}^m} &= \sum_j \theta_{xij}^m [r_{xij}^m - c^d d_{ij} - (w_{xi}^m + t_{xij})(c_x + \bar{u}) + V_{j|x_{im}} - V_{m|x_i}] \\ &\quad - \left| \frac{dw_{xi}(s_{xi}^m)}{ds_{xi}^m} \right| \sum_j q_{xij}^m (c_x + \bar{u} + V_{in|x_j} - V_{x+1j}) \end{aligned} \quad (3.35)$$

The intuition is simple. First, Equation (3.34) states that, to dispatch an additional passenger, platforms must compensate drivers for their opportunity cost of the trip, which is the sum of a cost of time and distance and a term $V_{m|x_i} - V_{j|x_{im}}$ trading off drivers' value at origin with their continuation value at destination.

As for Equation (3.35), the first term on the right-hand side captures captures a congestion effect. An additional idle driver reduces existing drivers' matching rates by $|d\theta_{xij}^m/ds_{xi}^m| = \theta_{xij}^m/s_{xi}^m$. When the number of drivers is increased, existing drivers must be compensated for the overall cost of this congestion externality. This cost is calculated as the number of existing drivers times the reduction in matching rates, $s_{xi}^m |d\theta_{xij}^m/ds_{xi}^m| = \theta_{xij}^m$, multiplied by the difference between the continuation values of being matched and staying idle, summed across all destinations. The second term

²⁷More precisely, the only remaining constraint is that allocations must be strictly positive. However, it can be shown that this constraint is never binding, provided that $\mathbf{p}_{xij}^m(q_{xij}^m, p_{xij}^{-m}, w_{xi}^m, w_{xi}^{-m}) \rightarrow \infty$ as $q_{xij}^m \rightarrow 0$, which holds for all common demand specifications. Hence, Problem 3.33 can be treated as unconstrained.

accounts for the fact that an additional idle driver reduces the wait times for all the $\sum_j q_{xij}^m$ trips originating from i by $|dw_{xi}^m/ds_{x,i}^m|$ hours. This results in a reduction in drivers' time-based costs of $c_x + \bar{u}$ dollars for every hour saved. Moreover, as trips become shorter, the likelihood that matched drivers reach destination before x transitions to x^{+1} increases. The quantity $V_{in|xj} - V_{x^{+1}j}$ measures the change in drivers' expected continuation value at destination for every hour saved.

The first order condition of Problem (3.33) with respect to q_{xij}^m yields

$$p_{xij}^m = \underbrace{c^d d_{ij} + (w_{xi}^m + t_{xij})(c_x + \bar{u})}_{\text{marginal cost of time and distance}} + \underbrace{\lambda^m q_{xij}^m \left| \frac{dp_{xij}^m}{dq_{xij}^m} \right|}_{\text{market power distortion}} + \underbrace{V_{m|xj} - V_{j|xim}}_{\text{continuation externality}}. \quad (3.36)$$

As captured by the first two terms, prices tend to be higher than the marginal cost of a trip, since platforms charge a mark-up proportional to their residual inverse demand elasticity. The last term is reminiscent of standard pricing formulae in two-sided markets: platforms charge a lower (resp. higher) price to users on one side of the market the higher is the positive (resp. negative) externality they create for users on the opposing side. Passengers create an externality on drivers by re-locating them from origin to destination, hence the cross-side externality trades off drivers' value at origin with their continuation value at destination. When $\lambda^m = 0$, the above equation yields the welfare-maximizing prices. If instead $\lambda^m > 0$, prices exceed the welfare-maximizing level, leading to an under-provision of passenger trips.

The first order condition of Problem (3.33) with respect to s_{xi}^m yields

$$\begin{aligned} & \underbrace{\sum_j \theta_{xij}^m [r_{xij}^m - c^d d_{ij} - (w_{xi}^m + t_{xij})(c_x + \bar{u}) + V_{j|xim} - V_{m|xj}]}_{\text{negative externality on idle drivers}} \\ &= \underbrace{\frac{dw_{xi}^m}{ds_{xi}^m} \frac{dW^m}{dw_{xi}^m}}_{\text{positive externality on riders}} + \underbrace{\lambda^m \frac{dw_{xi}^m}{ds_{xi}^m} \left[\sum_j q_{xij}^m \frac{dp_{xij}^m}{dw_{xi}^m} - \frac{dW^m}{dw_{xi}^m} \right]}_{\text{Spence distortion}} \quad (3.37) \\ &+ \underbrace{\left| \frac{dw_{xi}^m}{ds_{xi}^m} \right| \sum_j q_{xij}^m (c_x + \bar{u} + V_{in|xj} - V_{x^{+1}j})}_{\text{positive externality on matched drivers}}. \end{aligned}$$

Drivers are paid more (resp. less) the higher is the positive (resp. negative) externality they create for other users. Negative externalities stem from the congestion effect an additional driver creates on existing drivers by reducing their matching rates. Positive externalities stem from the fact that drivers reduce wait times, thus increasing passengers' surplus and leading to lower costs for matched drivers. When $\lambda^m = 0$, driver compensations are set to balance these opposing effects. If instead $\lambda^m > 0$, platforms do not fully internalize riders' surplus. Instead, they internalize that lower wait times allow for charging higher prices without losing passengers, hence drivers' compensations are distorted as in Spence (1975).

3.6. Gradient-based search. Proposition 6 is crucial from a computational standpoint, enabling the computation of equilibria through a gradient-based search for the allocations that maximize platforms' profits conditioned on the rival's strategy.

At every step, the search starts from a strictly positive allocation \mathbf{q}, \mathbf{s} and computes $\mathbf{w} = \mathbf{w}(\mathbf{s})$, $\boldsymbol{\theta} = \boldsymbol{\theta}(\mathbf{q}, \mathbf{s})$, and \mathbf{p} through passengers' inverse demand curves (Equation 3.2). The core of each step consists in inverting drivers' distribution, that is, computing $\boldsymbol{\delta} = \boldsymbol{\delta}(\mathbf{s}, \boldsymbol{\theta}, \mathbf{w})$ through the procedure outlined in Section 3.2.4.

Let dO^m/dq_{xij}^m and dO^m/ds_{xi}^m denote the derivatives of m 's objective evaluated at $\mathbf{q}^m, \mathbf{s}^m$, holding $\mathbf{p}^{-m}, \boldsymbol{\delta}^{-m}, \mathbf{w}^{-m}, \boldsymbol{\theta}^{-m}$ fixed. Once $\boldsymbol{\delta}$, \mathbf{w} and $\boldsymbol{\theta}$ are known, drivers' value functions can be computed via standard value iteration methods, hence these derivatives can be computed as per Proposition 6.²⁸

The search proceeds by updating \mathbf{q}, \mathbf{s} along these derivatives: given small a step size $\text{step} > 0$, the allocation at the next step, $\mathbf{q}_{(+1)}, \mathbf{s}_{(+1)}$, is computed as

$$q_{(+1)xij}^m = q_{xij}^m + \text{step} \cdot \frac{dO^m}{dq_{xij}^m}; \quad s_{(+1)xi}^m = s_{xi}^m + \text{step} \cdot \frac{dO^m}{ds_{xi}^m}.$$

The algorithm stops when $\mathbf{q}_{(+1)}, \mathbf{s}_{(+1)} \sim \mathbf{q}, \mathbf{s}$. In short, the algorithm simultaneously applies gradient ascent to both platforms' best response problems (Problem 3.33), and stops when the optimality conditions (Equations 3.36 and 3.37) are satisfied.

²⁸The derivative dO^m/dq_{xij}^m is simply the difference between the left- and right-hand sides of Equation (3.36). The derivative dO^m/ds_{xi}^m can be written as a function of $\boldsymbol{\delta}$ by taking the difference between the left- and right-hand sides of Equation (3.37), and substituting in Equation (3.28).

Convergence is not guaranteed with multiple platforms, although I consistently observe convergence across all simulations considered in Section 5. It is also worth noting that several counterfactuals involve a single platform, such as a monopolist or a social planner. With a single platform, a problem similar to 3.33 fully describes the model equilibria - rather than a platform's best response to the rival. Hence the gradient-based search reduces to standard gradient ascent, and convergence can be established under standard regularity conditions.

4. ESTIMATION

This section describes the estimation of the model primitives, namely, the demand systems, drivers' cost parameters, the matching technology, and platforms' conduct parameters. The Supplemental Material provides an evaluation of the model's fit.

4.1. **Demand.** The demand system is specified as follows. In hour x , M_{xij} potential passengers seek transportation from i to j , choosing between different ride-hailing platforms $m = 1, 2$ and an outside option. Their average utility from taking the outside option is normalized to zero, while the average utility from choosing Platform m is the sum of an average value v_{xij} , a disutility β_{xij}^p for each dollar spent, a disutility β_x^w per minute waited, and a brand effect α_{xi} if $m = 1$:

$$u_{xij}^m \equiv v_{xij} + \alpha_{xi} \cdot \{\text{if } m = 1\} - \beta_{xij}^p p_{xij}^m - \beta_x^w w_{xi}^m.$$

I assume that demand follows a nested logit model. The inner nest is a decision between platforms, conditional on taking a ride-hailing trip; the outer nest is a decision between ride-hailing and the outside option. This can be written as

$$\frac{q_{xij}^m}{M_{xij}} = \frac{\exp u_{xij}^m}{\sum_m \exp u_{xij}^m} \frac{(\sum_m \exp u_{xij}^m)^{\eta_{xij}}}{1 + (\sum_m \exp u_{xij}^m)^{\eta_{xij}}}$$

where η_{xij} is a nest parameter.

The parameters M_{xij} and η_{xij} are calibrated. M_{xij} is set to match the number of trips taken by New Yorkers across all transport modes from the 2018 MTA NYC travel survey.²⁹ η_{xij} is chosen so that the elasticity of total ride-hailing demand with respect

²⁹<https://new.mta.info/transparency/surveys>.

to prices is 0.6, in line with experimental estimates by Castillo (2023) and Cohen et al. (2016).³⁰

The price and wait time coefficients are estimated from high-frequency variation in platforms' relative shares. I divide the sample period in one-hour intervals t . Platform 1's relative share of passengers traveling from i to j at time t takes the logit formula

$$\frac{q_{tij}^1}{q_{tij}^1 + q_{tij}^2} = \frac{\exp(\alpha_{xi} - \beta_{xij}^p(p_{tij}^1 - p_{tij}^2) - \beta_x^w(w_{ti}^1 - w_{ti}^2))}{1 + \exp(\alpha_{xi} - \beta_{xij}^p(p_{tij}^1 - p_{tij}^2) - \beta_x^w(w_{ti}^1 - w_{ti}^2))} \quad (4.1)$$

where q_{tij}^m , p_{tij}^m and w_{ti}^m denote the number of trips, average passenger price and average wait time observed on Platform m at time t . I estimate this equation via maximum likelihood.

The idea behind identification is that, intuitively, demand shocks not affecting brand-specific preferences can impact the overall number of ride-hailing trips, but not platforms' relative shares. Hence endogeneity is a concern only insofar there is variation in brand preferences, which is controlled for by allowing the brand effects α_{xi} to vary granularly. In other words, identification relies on the assumption that, after controlling for weekday-hour- and location-specific fixed effects, the residual variation in relative prices and wait times is not due to idiosyncratic shocks to brand preferences - but rather to factors that can be treated as exogenous for demand estimation, such as shocks to driver availability of differences in platforms' pricing and matching algorithms. The maximum likelihood estimates of the brand effects, presented in the Supplemental Material, support this argument, as they suggest that brand preferences are fairly constant over time.

³⁰The elasticity of total ride-hailing demand with respect to prices is the percentage decrease in $\sum_m q_{xij}^m$ following a 1 percent increase in both p_{xij}^1 and p_{xij}^2 . This can be written as:

$$\sum_m \frac{q_{xij}^m}{\sum_m q_{xij}^m} \left[\frac{d \log q_{xij}^m}{d \log p_{xij}^m} + \frac{d \log q_{xij}^m}{d \log p_{xij}^{-m}} \right] = -0.6.$$

Once β_{xij}^p is determined, this equation yields a closed-form expression for η_{xij} .

Castillo (2023) studies an experiment where passengers in five Latin American cities were randomized into a control group and two treatment groups that received price discounts of 10% and 20%, finding an average demand elasticity of 0.633. Cohen et al. (2016) exploit quasi-experimental variation in Uber's pricing algorithm to estimate demand elasticities in several U.S. cities, finding an average elasticity of 0.6084 in NYC.

For estimation, I set $\beta_{xij}^p = \beta_i^p / t_{xij}$, capturing the fact that longer trips are less sensitive to absolute changes in prices.³¹ I let β_i^p vary by groups of regions and β_x^w vary across four-hour windows, capturing variation in income across locations and infra-day variation in passengers' value of time. α_{xi} varies across cells obtained from the combinations of regions, four-hour windows, and a weekday/weekend dummy. Table 2 below presents the maximum likelihood estimates of β_i^p and β_x^w , as well as the average estimates of passengers' value of time implied by these coefficients.³²

Region	β_i^p	VOT (\$/hr)	Hour of day	β_x^w	VOT (\$/hr)
Lower Manhattan	0.893 (0.011)	42.6	1am - 5am	0.015 (0.004)	12.6
Midtown	0.892 (0.010)	45.6	5am - 9am	0.026 (0.004)	25.2
Upper East and West	0.923 (0.015)	38.1	9am - 1pm	0.054 (0.005)	54
Harlem	0.972 (0.019)	31.3	1pm - 5pm	0.033 (0.003)	35.7
Bronx	0.991 (0.019)	31	5pm - 9pm	0.054 (0.002)	55.3
Brooklyn	0.882 (0.009)	39.3	9pm - 1am	0.029 (0.003)	26.9
Queens	0.944 (0.021)	33.1			

Table 2. Maximum likelihood estimates of passengers' sensitivity to price and wait time, represented by parameters β_i^p and β_x^w from Equation (4.1), along with the implied value of time estimates. The coefficients' standard errors are in parentheses.

The spatial variation in price coefficients is not large, ranging from 0.882 in Brooklyn to 0.991 in the Bronx. Wait time coefficients change substantially over time, as passengers are about 3.6 times more sensitive to wait times during the morning and

³¹This specification is consistent with the standard practice in the ride-hailing literature of estimating demand as a function of the surge multiplier, which assumes that passengers' price sensitivity is inversely proportional to a linear combination of time and distance (e.g. Cohen et al. (2016)).

³²Passengers' value of time is their willingness to pay to reduce wait time, as measured by the ratio $\beta_x^w / \beta_{xij}^p$. Table 2 presents region-specific and time-specific demand-weighted averages of these ratios.

evening rush hours than during late night. The average value of time varies both spatially and over time. It ranges from \$31 in the Bronx to \$45 in Midtown, and from \$12 at late night to \$55 during the evening rush hour. For comparison, estimates of hourly wages for 2019 Q3 assuming a 40-hour workweek are \$51 in Manhattan, \$24 in Brooklyn, \$27 in Queens and \$28 in the Bronx.³³ Another reference point is Buchholz et al. (2024), who present detailed estimates of the value of time for passengers of a ride-hailing platform in Prague. They find price elasticities four to ten times as large as wait time elasticities. Their findings are roughly consistent with my estimates, which imply price elasticities five to twenty times as large as wait time elasticities.

4.2. Cost parameters and labor supply elasticities. Recall that drivers’ cost parameters include the per-mile cost c_d , the time-varying outside options c_x , a parameter γ capturing heterogeneity in drivers’ outside options, and a parameter \bar{u} capturing registration costs.

γ and c_d are calibrated. γ is set so that drivers’ labor supply elasticity on the intensive margin is 0.5, consistently with experimental estimates by Hall et al. (2023) and Caldwell and Oehlsen (2022).³⁴ c_d is set equal to the July 2019 standard rate of \$1.088 per mile set by the TLC.³⁵

The time-varying outside options c_x are estimated from temporal and spatial variation in drivers’ wages and labor supply. Estimation relies on the identification result and

³³According to the Quarterly Census of Employment and Wages by the U.S. Bureau of Labor Statistics, the average weekly wage in 2019 Q3 was \$2,055 in New York County (Manhattan), \$955 in Kings County (Brooklyn), \$1,074 in Queens County, and \$1,130 in Bronx County.

³⁴Drivers’ labor supply elasticity on the intensive margin is the percentage increase in $\sum_x n_x$, drivers’ total hours worked per week, following a percentage increase in drivers’ hourly wages, holding fixed N , the number of drivers in the NYC labor pool. Drivers’ hourly wages are the sum of total weekly payments to drivers, $\sum_{xijm} q_{xij}^m r_{xij}^m$, divided by the total hours worked per week. This elasticity does not admit a closed-form expression, so I set γ so that hours worked increase by 5% following a 10% increase in hourly wages.

Hall et al. (2023) study drivers’ reactions to an exogenous 40% increase in hourly earnings across several US cities, finding that hours worked per driver increased by 20%. Caldwell and Oehlsen (2022) study an experiment where drivers in Houston were randomized into a control group and a treatment group that received 10-50% higher earnings per trip, finding an intensive margin labor supply elasticity of 0.8 for female drivers and 0.4 for male drivers. According to Cook et al. (2021), 27% of drivers in the US between 2015 and 2017 were women, resulting in an average elasticity of 0.508.

³⁵This accounts for expenses such as “vehicle purchase or lease, fuel, maintenance, and insurance”: www.nyc.gov/assets/tlc/downloads/pdf/driver_income_rules_12_04_2018.pdf

inversion algorithm outlined in Section 3.2.4: once the number of drivers N and their distribution s are observed, their expected payoffs δ_{xi}^m from one hour spent searching at various locations are identified and can be computed. These payoffs were defined in Equation (3.28), repeated here for simplicity:

$$\delta_{xi}^m = -c_x + \sum_j \theta_{xij}^m [r_{xij}^m - c_d d_{ij} - (w_{xi}^m + t_{xij})c_x]. \quad (4.2)$$

Intuitively, for each hour spent searching, drivers incur the outside option c_x and get dispatched $\sum_j \theta_{xij}^m$ times on average, receiving the compensation net of the cost of time and distance. To compute δ , N is set equal to 80000, which is the number of unique drivers who dispatched at least one ride-hailing trip in July 2019 according to TLC data. Once δ and c_d are known, Equation (4.2) yields a system of equations whose only unknowns are the time-varying parameters c_x . Figure 4.1 plots their least squares estimates.




Figure 4.1. Drivers' time-varying outside options c_x , measured in dollars per hour. Dashed lines indicate the 95% confidence intervals.

Drivers' outside options are approximately \$35 per hour on average. The estimated value ranges from around \$33 per hour during the day to roughly \$39 per hour during late night, indicating that leisure utility may be greatest during nighttime hours. A potential point of comparison is the hourly wage earned by construction workers in the state of New York, which averaged roughly \$38 in June 2019.³⁶

Once c_x and c_d are known, solving drivers' dynamic choice problem yields their average weekly revenues net of outside options u . The estimated value of u is very

³⁶See <https://labor.ny.gov/stats/ceshourearn2.asp>

small, at about \$3660 per year. For comparison, direct costs that drivers must pay for licensing related requirements total slightly over \$1500.³⁷ This is because drivers in the sample work an average of only 2 hours per day. On the one hand, few hours worked per week translate directly into low net weekly revenues. On the other hand, the model rationalizes the small number of hours worked with high outside options, resulting in small net revenues per hour worked. This fact likely reflects the real situation of NYC drivers in 2019: after accounting for calibrated vehicle-related expenses, drivers make about \$15 per hour on average. This is the NYC minimum wage in 2019, which is a good estimate of drivers’ outside options per hour worked. Some authors argue that NYC drivers make even below this threshold (Parrott and Reich, 2018).

Note that, multiplying u by the total driver labor force, which is 80,000 drivers in July 2019, yields only \$293 million per year as an upper bound for the total drivers’ surplus at status quo that would obtain if free entry (Equation 3.7) was replaced with any extensive margin labor supply curve. For comparison, riders make about \$4,097 million per year in surplus, and platforms make about \$669 million per year in profits. Intuitively, this makes the market empirically close to one with drivers’ free entry on the extensive margin, justifying Equation (3.7). If free entry was relaxed in counterfactuals, drivers’ surplus would change little in absolute terms relative to other quantities. Specifically, changes in drivers’ surplus would not significantly impact total welfare calculations, which are dominated by changes in riders’ surplus. The Supplemental Material presents a robustness analysis that confirms this intuition across alternative calibrations.

4.3. Matching function. A known feature of ride-hailing markets is that observed wait times tend to be well-approximated by a constant elasticity function of the spatial density of idle drivers.³⁸ I observe the same in my sample when estimating a

³⁷The TLC itemizes these costs in a flier distributed to prospective drivers:

http://www.nyc.gov/html/tlc/downloads/pdf/thinking_about_driving_fhv.pdf

³⁸See Yan et al. (2020) and Castillo et al. (2023). Yan et al. (2020, Proposition 1) show that this specification can also be justified theoretically: a simple model of matching in a homogeneous, two-dimensional space predicts a wait time that is a constant elasticity function of the density of idle drivers. I thank an anonymous referee for highlighting this.

constant elasticity function of the form

$$\log w_{xi}^m = \alpha_i - \beta^{\text{speed}} \log(\text{speed}_{xi}) - \beta^{\text{density}} \log(\text{driver density}_{xi}^m) \quad (4.3)$$

where α and β are the parameters to be estimated. speed_{xi} denotes the average traffic speed in region i in weekday-hour x , while $\text{driver density}_{xi}^m$ denotes the ratio between the number of idle drivers s_{xi}^m and the area of region i . Intuitively, higher driver density implies that passengers are matched to closer drivers, thus wait times are lower. This relationship is affected by the speed of traffic and by topological features, e.g. road patterns, which are controlled for by the region-specific fixed effects α_i .

log(pick-up time)			
log(speed)	-1.554 (0.011)	R ²	0.809
log(driver density)	-0.666 (0.003)	Obs	6552
Region fixed effects	Yes		

Table 3. Least squares estimates of the matching function coefficients.

Table 3 shows the least squares estimates of Equation (4.3). If the traffic speeds double, wait times are reduced by 150%. The role driver density is also substantial, as a one-fold increase in density reduces wait times by about 55%. While parsimonious, this model explains a substantial share of the space-time variation in wait times, as measured by a coefficient R^2 of 0.8.

4.4. Conduct parameters. Using the estimated demand parameters, cost parameters, and matching technology, the optimal pricing formulae in Equations (3.36) and (3.37) yield a closed-form expression of the optimal markups in prices (i.e. the difference between passenger fares and driver earnings, net of taxes and tolls) as a function of platforms' conduct parameters. I calibrate the conduct parameters to match the average markups observed in the data, which leads to $\lambda^1 = 0.221$ and $\lambda^2 = 0.305$. This means that platforms charge passengers below the profit maximization level, consistently with analogous findings in ride-hailing (Castillo, 2023) and similar industries (e.g. Gutierrez, 2021 and Sullivan, 2022).

5. COUNTERFACTUALS

I simulate counterfactuals varying two main primitives.³⁹ First, I consider counterfactuals where a single entity - a monopolist, or a social planner - sets prices on both platforms to maximize the joint objective

$$O^* = \lambda \sum_m \Pi^m + (1 - \lambda) \sum_m (\Pi^m + CS^m). \quad (5.1)$$

When $\lambda > 0$, this leads to a monopolist maximizing a combination of profits and total surplus across both marketplaces. When $\lambda = 0$ it leads to a social planner maximizing total welfare.

Second, I consider counterfactuals where the two platforms are interoperable. Recall that active drivers divide their time between idling and dispatching trips. Importantly, idle drivers may be available on either platform, but not on both. This is captured by Equation (3.6), repeated here for simplicity:

$$\sum_{im} s_{xi}^m + \sum_{ijm} q_{xij}^m (w_{xi}^m + t_{xij}) = n_x. \quad (5.2)$$

I compare this with counterfactuals where idle drivers are simultaneously available on both platforms, hence the above equation is replaced with

$$s_{xi}^1 = s_{xi}^2 \equiv s_{xi}; \quad \sum_i s_{xi} + \sum_{ijm} q_{xij}^m (w_{xi}^m + t_{xij}) = n_x. \quad (5.3)$$

Comparing Equations (5.2) and (5.3), holding demand q_x and drivers' labor supply n_x fixed, the number of drivers available to m 's passengers, $\sum_i s_{xi}^m$, is always larger when platforms are interoperable. In the status quo, passengers may miss potential matches with nearby drivers who are active on a different platform, leading to longer wait times (as per Equation 3.8). Interoperability eliminates this inefficiency by making all idle drivers simultaneously available to passengers on both platforms.

5.1. Efficiency. I start by computing the social planner solution in two alternative scenarios: one in which the two platforms are interoperable, and one in which they are not. In both scenarios, the social planner sets prices on both platforms to maximize

³⁹In these counterfactuals, the supply-side dynamics described in Section 3.2 and platforms' first order conditions (Equations 3.36 and 3.37) change slightly. These changes are detailed in the Supplemental Material (Rosaia 2024).

total welfare, under the constraint that drivers' labor supply should not exceed an upper bound, i.e. $\sum_x n_x \leq \bar{n}$. I compute a solution to this problem for different values of \bar{n} , tracing an efficient frontier that measures the maximum welfare achievable in equilibrium as a function of total amount of hours drivers spend on the street.

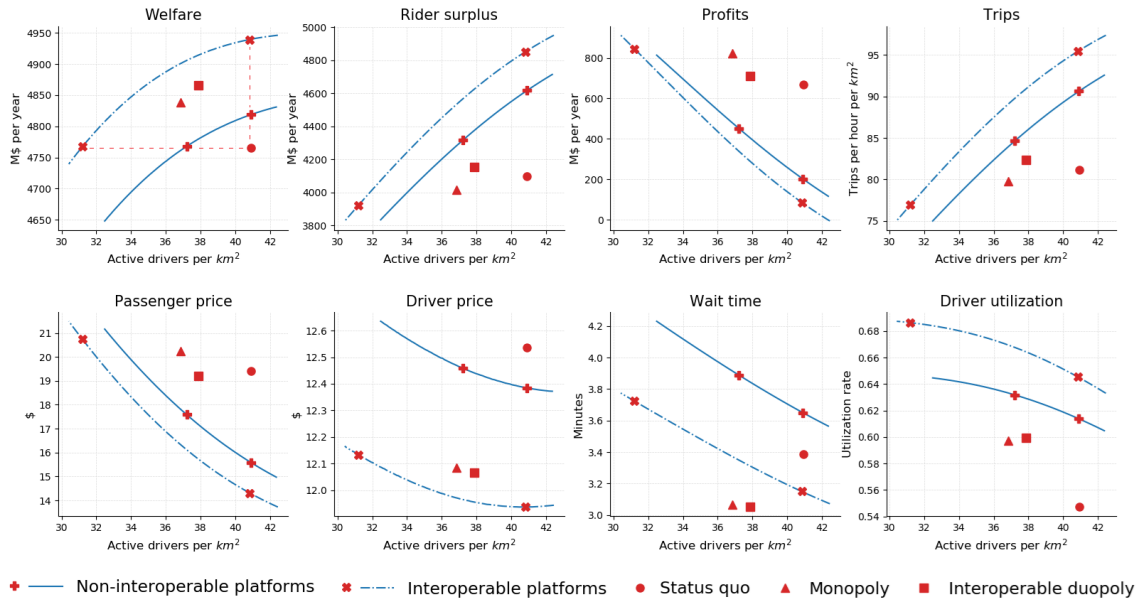


Figure 5.1. Efficient frontiers with and without interoperability, welfare-maximizing equilibria, traffic-minimizing equilibria, and equilibria under monopoly and interoperable duopoly.

The solid lines in Figure 5.1 depict the efficient frontier when platforms are not interoperable, with drivers' average density on the x-axis. Table 4 shows the changes in market outcomes corresponding to two points along this frontier.

First, the intersection with the vertical line originating from the status quo in the first panel defines a welfare-maximizing equilibrium. The vertical distance between this point and the status quo represents an absolute loss: if platforms were to set prices efficiently, welfare would increase by \$53 million per year without adding more drivers to the streets, that is, without worsening the traffic externalities imposed by ride-hailing services. Intuitively, this is due to market power causing passenger prices to be set 20% above the constrained welfare-maximizing level. This leads to sub-optimal trip demand, with platforms dispatching 12% fewer trips. Consequently,

drivers receive fewer requests, hence they spend more time idle, and their utilization is sub-optimal. In essence, drivers are not utilized efficiently: by keeping drivers busier, a social planner could dispatch 12% more trips, increasing rider surplus by \$520 million per year without affecting overall traffic. Note that drivers would earn 1% less per trip, but this would be offset by dispatching more trips. However, the reduction in prices would outweigh the reduction in costs, leading to a 70% drop in profits, which is not aligned with platforms' private incentives.

Second, the intersection with the horizontal line defines a traffic-minimizing equilibrium. The horizontal distance between this point and the status quo represents an absolute traffic waste: traffic could be reduced by 9% without reducing welfare. The rationale is similar to the previous scenario. By keeping drivers busier, a social planner could dispatch 4% more trips and increase rider surplus by 5%, all while reducing overall traffic by 9%. However, the slight decrease in costs per trip (-0.6%) would not offset the significant price reduction (-9%) needed to boost demand, and profits would drop by 32%. Even though the increase in rider surplus exactly offsets the profit reduction from a social welfare perspective, this configuration does not align with platforms' private incentives.

The dashed lines depict the efficient frontier when platforms are interoperable. In the first panel, the efficient frontier is shifted upwards. This is because interoperability improves the efficiency of the matching process. This leads to shorter wait times and higher driver utilization, resulting in lower costs per trip, hence lower prices, more trips and higher surplus for riders. The intersections with the vertical line and the horizontal line define new welfare-maximizing and traffic-minimizing equilibria. The vertical distance between the two lines measures the additional welfare wasted due the lack of interoperability between the two platforms: if platforms were interoperable, welfare could be increased by an additional \$120 million per year without adding more drivers to the streets. The horizontal distance measures the additional traffic wasted: traffic could be reduced by an additional 15% without reducing welfare.

5.2. Market structure. This section explores the extent to which these inefficiencies can be mitigated under alternative market structures. First, I simulate a scenario where a monopolist sets prices on both platforms to maximize the joint objective stated in Equation (5.1). The conduct parameter is set equal to the value estimated for

			Non-interoperable platforms		Interoperable platforms	
	Monopoly	Interoperable duopoly	Welfare maximizing	Traffic minimizing	Welfare maximizing	Traffic minimizing
Welfare (M\$/yr)	+1.5% (+72.2)	+2.1% (+99.7)	+1.1% (+52.9)	- -	+3.6% (+173.5)	- -
Rider surplus (M\$/yr)	-2% (-82.4)	+1.4% (+56.4)	+12.7% (+519.6)	+5.3% (+218.7)	+18.4% (+756.1)	-4.3% (-175.5)
Profits (M\$/yr)	+23.1% (+154.6)	+6.5% (+43.3)	-69.9% (-466.7)	-32.4% (-216.6)	-87.1% (-582.6)	+26.5% (+177.1)
Trips	-1.7%	+1.5%	+11.7%	+4.3%	+17.7%	-5.2%
Active drivers	-10%	-7.5%	-	-9.1%	-	-23.7%
Driver utilization	+9.2%	+9.5%	+12.2%	+5.5%	+18%	+25.4%
Passenger prices	+4.2%	-1.2%	-19.8%	-9.5%	-26.3%	+6.8%
Driver prices	-3.6%	-3.8%	-1.2%	-0.6%	-4.8%	-3.3%
Wait times	-9.4%	-9.9%	+7.8%	+14.9%	-7%	+10%

Table 4. Counterfactual market outcomes measured in changes relative to the status quo. The welfare-maximizing and traffic-minimizing equilibria are defined as the points where the efficient frontiers intersect with the vertical and horizontal lines originating from the status quo.

Platform 1 - i.e. $\lambda = \lambda^1 = 0.221$. To abstract from changes in welfare due to changes in service variety, passengers keep the ability to choose between the two separate platforms, hence their demand curves (Equation 3.1) are unchanged. However, the two platforms are interoperable, hence the specific platform chosen by passengers is irrelevant from the monopolist's perspective. This scenario can be interpreted as a merger, or more specifically as the outcome of Platform 1 acquiring Platform 2.

Making platforms interoperable allows the monopolist to achieve a more efficient matching process. However, eliminating competition leads the monopolist to raise prices, worsening the market power distortions and offsetting most efficiency gains. The new equilibrium is depicted in Figure 5.1, while Table 4 displays the changes in market outcomes relative to the status quo, and Figure 5.2 illustrates the breakdown of these changes across regions.

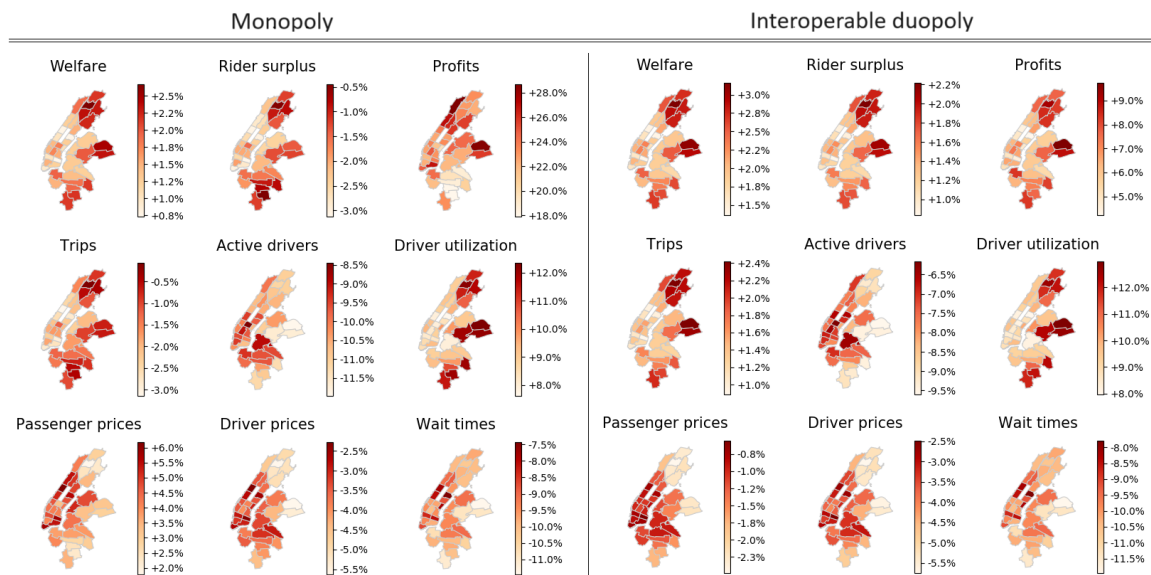


Figure 5.2. Market outcomes under a monopoly and an interoperable duopoly, measured in percentage changes relative to the status quo.

Compared to the status quo, the merger improves matching efficiency, reducing wait times by 9% and increasing driver utilization by the same margin. By keeping drivers busier, the monopolist reduces traffic by 10% while reducing passenger trips by only 2%. Higher driver utilization also leads to cost savings per trip, allowing the monopolist to reduce driver prices by 3.6%. The overall impact on welfare is positive. However, these efficiency gains are not passed on to consumers, who instead face a 4% price increase, resulting in a \$82 million annual drop in their surplus. Most efficiency gains are captured by the monopolist, whose profits increase by \$155 million per year. These figures vary significantly across space, with low-density peripheral regions experiencing the highest efficiency gains. Intuitively, matching frictions are more pronounced in low-density areas, where passengers often get matched with far-away drivers. Increasing density in these areas leads to substantial reductions in wait times. As drivers spend less time to pick up passengers, their utilization increases, leading to large cost savings. A portion of these cost savings are passed on to riders, who experience a smaller increase in prices and a smaller decrease in ridership and surplus following a merger. Conversely, matching frictions are minimal in the high-density regions of southern Manhattan, where riders and drivers are in close proximity and

wait times are already near their minimum. Increasing density in these areas results in smaller reductions in wait times, a smaller increase in driver utilization, and lower cost savings, leading to higher price increases and a greater decrease in ridership and rider surplus following a merger.

To summarize, the merger achieves both positive and negative outcomes. By making platforms interoperable, it reduces traffic without reducing overall welfare. However, the increased market power adversely affects riders, especially in high-density regions. This raises the question of whether a market configuration exists that achieves the benefits of interoperability without losing the advantages of competition.

In the last counterfactual, I simulate a scenario where two interoperable platforms compete with each other. This means that, on the one hand, platforms independently set prices to maximize their own objectives. On the other hand, at any given time, all idle drivers are simultaneously available on both platforms. In practice, this can be achieved by removing the obstacles and incentives preventing drivers to simultaneously multi-home, as discussed in Section 2.3.

As shown in Figures 5.1 and 5.2 and Table 4, in many respects an interoperable duopoly achieves outcomes comparable to those of a merger: it reduces wait times and keeps drivers busier, lowering costs and reducing traffic. The key difference is that a significant portion of these efficiency gains are now passed on to riders, who benefit from lower prices and increased surplus, and take more trips. This is true across the entire city, with passengers in peripheral areas benefiting the most. Interestingly, platforms' profits also rise, albeit by a considerably smaller margin compared to the merger scenario.

5.3. The role of conduct. These results are specific to the conduct observed in the sample: if conduct changes, equilibria would also change, and there is no easy way to predict how conduct might change over time. Conduct can be endogenous or reflect exogenous strategic decisions made at the firm or city level, or be influenced by the ownership structure.

Since these dynamics cannot be assessed with the data at hand, the Supplemental Material includes a robustness analysis to explore how counterfactual outcomes change with conduct. This provides both reassurance and caution. On one hand,

most qualitative insights remain robust to exogenous changes in conduct. In terms of magnitudes, the results are stable for the interoperable duopoly but vary for the merger: as λ increases, a monopoly leads to higher price increases, resulting in larger drops in ridership and surplus. Specifically, this means that the total welfare calculations may flip sign. With values of λ slightly higher than those estimated, the increase in profits no longer compensates for the decrease in rider surplus, leading to a reduction in total welfare.

Therefore, the conclusion that a monopoly might increase total welfare is highly context-dependent and generally does not hold. The welfare reduction reaches a maximum of \$295 million per year when $\lambda = 1$ (i.e., when platforms maximize plain profits).

6. CONCLUSION

The lack of integration between ride-hailing networks results in sub-optimal outcomes such as higher prices, reduced mobility, and increased traffic. However, consolidating users into a single network involves a tension between higher efficiency and increased market power. This paper provides a framework to quantify these opposing forces. A merger, while potentially improving efficiency, would adversely affect riders, especially in high-density markets. In contrast, removing barriers to simultaneous multi-homing would capture the benefits of integration without losing the advantages of competition, reducing traffic and improving mobility while increasing profits and consumer welfare.

REFERENCES

- AHLFELDT, G. M., S. J. REDDING, D. M. STURM, AND N. WOLF (2015): “The economics of density: Evidence from the Berlin Wall,” *Econometrica*, 83, 2127–2189.
- ALLON, G., D. CHEN, AND K. MOON (2023): “Measuring strategic behavior by gig economy workers: Multihoming and repositioning,” *Working Paper*.
- ALMAGRO, M., F. BARBIERI, J. C. CASTILLO, N. G. HICKOK, AND T. SALZ (2024): “Optimal Urban Transportation Policy: Evidence from Chicago,” *Working Paper*.
- ARGENTESI, E. AND L. FILISTRUCCHI (2007): “Estimating market power in a two-sided market: The case of newspapers,” *Journal of Applied Econometrics*, 22, 1247–1266.

- ARMSTRONG, M. (2006): “Competition in two-sided markets,” *The RAND Journal of Economics*, 37, 668–691.
- BESBES, O., F. CASTRO, AND I. LOBEL (2021): “Surge pricing and its spatial supply response,” *Management Science*, 67, 1350–1367.
- BESEN, S. M. AND L. L. JOHNSON (1986): *Compatibility standards, competition, and innovation in the broadcasting industry*, Rand Corporation.
- BIAN, B. (2020): “Search Frictions, Network Effects and Spatial Competition: Taxis versus Uber,” *Working paper*.
- BRANCACCIO, G., M. KALOUPTSIDI, AND T. PAPAGEORGIU (2020): “Geography, Transportation, and Endogenous Trade Costs,” *Econometrica*, 88, 657–691.
- BRANCACCIO, G., M. KALOUPTSIDI, T. PAPAGEORGIU, AND N. ROSAIA (2023): “Search frictions and efficiency in decentralized transport markets,” *The Quarterly Journal of Economics*, 138, 2451–2503.
- BRESNAHAN, T. (2002): “The economics of the Microsoft case,” *Working paper*.
- BROCK, G. W. (1981): “The telecommunications industry: The dynamics of market structure,” *Harvard University Press*.
- BUCHHOLZ, N. (2022): “Spatial equilibrium, search frictions, and dynamic efficiency in the taxi industry,” *The Review of Economic Studies*, 89, 556–591.
- BUCHHOLZ, N., L. DOVAL, J. KASTL, F. MATEJKA, AND T. SALZ (2024): “Personalized Pricing and the Value of Time: Evidence from Auctioned Cab Rides,” *Working paper*.
- CALDWELL, S. AND E. OEHLSEN (2022): “Gender, outside options, and labor supply: Experimental evidence from the gig economy,” *Working Paper*.
- CAO, G., G. Z. JIN, X. WENG, AND L.-A. ZHOU (2021): “Market-expanding or Market-stealing? Competition with network effects in bike-sharing,” *The RAND Journal of Economics*, 52, 778–814.
- CASTILLO, J. C. (2023): “Who Benefits from Surge Pricing?” *Econometrica*, Forthcoming.
- CASTILLO, J. C., D. T. KNOEPFLE, AND E. G. WEYL (2023): “Matching and pricing in ride hailing: Wild goose chases and how to solve them,” *Management Science*, Forthcoming.
- CHANDRA, A. AND A. COLLARD-WEXLER (2009): “Mergers in two-sided markets: An application to the Canadian newspaper industry,” *Journal of Economics & Management Strategy*, 18, 1045–1070.
- CLEMENTS, M. T. AND H. OHASHI (2005): “Indirect network effects and the product cycle: video games in the US, 1994–2002,” *The Journal of Industrial Economics*, 53, 515–542.

- COHEN, P., R. HAHN, J. HALL, S. LEVITT, AND R. METCALFE (2016): “Using big data to estimate consumer surplus: The case of uber,” *Working Paper*.
- COOK, C., R. DIAMOND, J. V. HALL, J. A. LIST, AND P. OYER (2021): “The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers,” *The Review of Economic Studies*, 88, 2210–2238.
- CORTS, K. S. AND M. LEDERMAN (2009): “Software exclusivity and the scope of indirect network effects in the US home video game market,” *international Journal of industrial Organization*, 27, 121–136.
- EVANS, D. AND R. SCHMALENSEE (2005): “The economics of interchange fees and their regulation: An overview,” .
- FARRELL, J. AND P. KLEMPERER (2007): “Coordination and lock-in: Competition with switching costs and network effects,” *Handbook of industrial organization*, 3, 1967–2072.
- FARRELL, J., C. SHAPIRO, R. R. NELSON, AND R. G. NOLL (1992): “Standard setting in high-definition television,” *Brookings Papers on Economic Activity. Microeconomics*, 1992, 1–93.
- FRECHETTE, G. R., A. LIZZERI, AND T. SALZ (2019): “Frictions in a Competitive, Regulated Market Evidence from Taxis,” *American Economic Review*, 109, 2954–2992.
- GABEL, H. L. (1991): “Competitive strategies for product standards: The strategic use of compatibility standards for competitive advantage,” *McGraw-Hill*.
- GAINEDDENOVA, R. (2022): “Pricing and efficiency in a decentralized ride-hailing platform,” *Working paper*.
- GENAKOS, C. AND T. VALLETTI (2011): “Testing the "waterbed" effect in mobile telephony,” *Journal of the European Economic Association*, 9, 1114–1142.
- GHILI, S. AND V. KUMAR (2021): “Spatial distribution of supply and the role of market thickness: Theory and evidence from ride sharing,” *arXiv preprint arXiv:2108.05954*.
- GUTIERREZ, G. (2021): “The welfare consequences of regulating Amazon,” *Working paper*.
- HALL, J. V., J. J. HORTON, AND D. T. KNOEPFLE (2023): “Ride-sharing markets re-equilibrate,” *Working paper*.
- JEZIORSKI, P. (2014): “Effects of mergers in two-sided markets: The US radio industry,” *American Economic Journal: Microeconomics*, 6, 35–73.
- JULLIEN, B., A. PAVAN, AND M. RYSMAN (2021): “Two-sided markets, pricing, and network effects,” *Handbook of Industrial Organization*, 4, 485–592.
- KAISER, U. AND J. WRIGHT (2006): “Price structure in two-sided markets: Evidence from the magazine industry,” *International Journal of Industrial Organization*, 24, 1–28.

- KATZ, M. L. AND C. SHAPIRO (1985): “Network externalities, competition, and compatibility,” *The American economic review*, 75, 424–440.
- KREINDLER, G. (2023): “Peak-hour road congestion pricing: Experimental evidence and equilibrium implications,” *working paper*.
- KREINDLER, G., A. GADUH, T. GRAFF, R. HANNA, AND B. A. OLKEN (2023): “Optimal Public Transportation Networks: Evidence from the World’s Largest Bus Rapid Transit System in Jakarta,” .
- LEE, R. S. (2013): “Vertical integration and exclusivity in platform and two-sided markets,” *American Economic Review*, 103, 2960–3000.
- LEHR, W. (1995): “Compatibility standards and interoperability: Lessons from the Internet,” *Standards policy for information infrastructure*, 121.
- LIU, T. X., Z. WAN, AND C. YANG (2021): “The Efficiency of A Dynamic Decentralized Two-sided Matching Market,” *Working Paper*.
- MA, H., F. FANG, AND D. C. PARKES (2022): “Spatio-temporal pricing for ridesharing platforms,” *Operations Research*, 70, 1025–1041.
- MAGNAC, T. AND D. THESMAR (2002): “Identifying dynamic discrete decision processes,” *Econometrica*, 70, 801–816.
- OSTROVSKY, M. AND M. SCHWARZ (2018): “Carpooling and the Economics of Self-Driving Cars,” *Mimeo, Stanford University*, 25487.
- PARROTT, J. A. AND M. REICH (2018): “An Earnings Standard for New York City’s App-based Drivers: Economic Analysis and Policy Assessment,” .
- ROCHET, J.-C. AND J. TIROLE (2003): “Platform competition in two-sided markets,” *Journal of the european economic association*, 1, 990–1029.
- ROSAIA, N. (2024): “Supplement to ‘Competing Platforms and Transport Equilibrium’,” *Working Paper*.
- RYSMAN, M. (2004): “Competition between networks: A study of the market for yellow pages,” *The Review of Economic Studies*, 71, 483–512.
- SCHALLER, B. (2017): “Empty Seats, Full Streets. Fixing Manhattan’s Traffic Problem,” .
- SHAPIRO, M. H. (2018): “Density of Demand and the Benefit of Uber,” *Working paper*.
- SPENCE, A. M. (1975): “Monopoly, quality, and regulation,” *The Bell Journal of Economics*, 417–429.
- SULLIVAN, M. (2022): “Price controls in a multi-sided market,” *Working paper*.
- YAN, C., H. ZHU, N. KOROLKO, AND D. WOODARD (2020): “Dynamic pricing and matching in ride-hailing platforms,” *Naval Research Logistics (NRL)*, 67, 705–724.

YANG, R. (2022): “(Don’t) Take Me Home: Home Preference and the Effect of Self-Driving Trucks on Interstate Trade,” *Working paper*.