Multi-homing on Heterogeneous Taxi Hail Modes: Welfare Impact of Uber's partnership with Yellow Taxi companies in NYC

By SAHIL KHATKAR *

In March 2022, Uber announced that it would allow yellow taxi drivers to join the app in NYC and subsequently began rolling out the feature in September 2022. This was the first such large-scale agreement of its kind in the US. Under the agreement, yellow taxi drivers can choose when to be available on the app. Using publicly available pre-announcement ride data from NYC, I estimate a nested logit demand model and compute a dynamic supply algorithm which provides vacant cab distributions for yellow taxi and Ubers with different matching technologies. By computing a counterfactual which allows yellow taxi drivers to choose both location and hail mode while searching, I find that this additional choice raises aggregate consumer surplus 1.89 times, increases the profits for those using Uber hail mode by reducing wait times due to increased supply, and reduces the profits for those yellow taxi who continue to street hail.
JEL: C78, D83, L91, R12, R41
Keywords: Transportation, Search, Spatial Equilibrium

Before the Covid pandemic, ridehailing apps in NYC were often considered the bane of the yellow taxi industry. While market shares of ride-hailing apps (in terms of number of rides) increased from 12% in Jan 2015 to 84% in July 2022 (quoted from TLC Trip Record data), the value of yellow taxi medallions plummeted from $1 million in 2014 to $100,000 in 2021. This led to significant medallion debt for the yellow taxi (henceforth YT) drivers (to the tune of $600,000 for an average medallion owner[1]) even leading NYC to announce a substantive Medallion Relief Program.

During the pandemic, however, both the demand and supply of cabs fell off a cliff, with aggregate trips per day across platforms reducing by 84% from 966,000 in Feb, 2020 to 151,000 in April, 2020. As travel restrictions eased over time, ride-hailing companies claimed that not enough drivers were returning to the platform to meet the rising demand. In a May 2021 earnings report, Uber reported active drivers were down 22% year-on-year.

In response, Uber announced its first large scale partnership of the kind (within

* Contact: The University of Texas at Austin, Economics Department, sahil.khatkar@utexas.edu.
[1]Source: Cecilia Saixue Watt (Oct 2017), "There's no future for taxis: New York yellow cab drivers drowning in debt," *The Guardian*

the US) with YT companies in NYC in the last week of March, 2022.[2] Under this agreement, all of NYC's YT would appear on the Uber app. The YT rides booked through the app would be priced similar to Uber's standard option, UberX, rather than the metered fare from street-hail and Uber would get a share of the fare amount. Both the passengers and drivers would also be able to see the fares upfront.

On the face of it, one could argue that this is a perfect match between a platform looking for drivers and drivers looking for rides. However, the overall welfare impact of such a change are not so straightforward, nor is it easy to find who the welfare would accrue to. For instance, a significant increase in drivers on the Uber app could potentially reduce prices and lower aggregate revenue for each driver. On the other hand, if enough YT drivers switch to the app, the existing street-hailing market's thickness could reduce to the point where demand reduces because expected wait times are too high. This would be detrimental to any YT driver who doesn't make the switch.

These impacts could, of course, be potentially offset by the aggregate welfare gains associated with the various technological advantages that Uber offers. Existing literature that I discuss below shows, however, that these welfare gains depend heavily on the density of the market (defined as the number of potential passengers). Therefore, an analysis of the agreement's welfare impact and its distribution needs to take into account this spatial variation.

For the case of NYC, the massive debt accrued by YT drivers may provide some rationale for why such an analysis is economically relevant to the stakeholders, but Uber has also been expanding to other cities (e.g., San Francisco[3]). Therefore, such an analysis could potentially have wider applications than just this particular case.

I attempt to at least partially analyze the welfare impacts of this agreement by using a counterfactual analysis where YT, when not employed on a trip, can choose where to be and what hail mode to use, street-hail or the Uber app. Using a nested logit demand and a dynamic spatial supply algorithm, I compute welfare impacts where I incorporate some potential consequences like the response of Uber drivers to additional yellow taxis and the change in demand due to wait times.

Briefly, the nested logit is estimated using weekday 6 AM to 4 PM observed ride level data and approximate public transport shares for a subset of zones (at the border of Manhattan and Queens/Brooklyn) in NYC for March 2022. The supply algorithm follows Buchholz (2022) and Shapiro (2020) and uses a finite horizon dynamic framework where unmatched cabs make spatial choices based on policy functions that use logit shares to define travel probabilities. These shares are dependent on expectations over continuation values, which incorporate the probability of matching with a passenger and the expected profit from the

[2]Source: Hu, W., Browning, K. and Zraick, K. (March 29, 2022), "Uber Partners With Yellow Taxi Companies in N.Y.C.," *The New York Times*

[3]Source: White, M. (2022), "Once seemingly impossible, Uber partnership with S.F. taxis begins," San Francisco Examiner

match. Uber and YT's match probabilities are modelled differently to reflect their different technologies.

The counterfactual allows YT to make simultaneous location and hail mode choices. Ubers re-optimize their travel patterns when YT drivers use the app, and wait times implied by the model are used to update demand. The algorithm is computed till convergence. Using this, I find that, on average in a day, around 15% YT who were searching in Manhattan switch to the app and shift to the boroughs. The cabs who remain in Manhattan are almost evenly split between the two hail modes. On the other hand, Ubers' change in search behavior is localized to the regions.

Consumer surplus triples in the boroughs as more YT (on Uber) enter the market and reduce app-hail wait times, whereas it remains the same in Manhattan. The YT who continue to street-hail in Manhattan suffer losses in revenue, whereas those who switch to the app have increased profits in both the regions.

This counterfactual doesn't, however, capture two salient features which would impact the results. a) I do not model Uber's price setting mechanism and keep prices fixed in the counterfactual, and b) due to data constraints, I'm unable to incorporate entry/exit decisions in the model. These are major limitations that I hope to resolve in the future by changing the model specification to allow price setting, and using a different data set which identifies entry/exit patterns.

Related Literature Broadly, this paper incorporates ideas from the search and matching dynamic spatial equilibrium literature dealing with urban transportation. Specifically, the supply side computation follows the methods used by Buchholz (2022) and Shapiro (2020), while the demand side is a standard nested logit.

Buchholz (2022) computes the dynamic spatial equilibrium model[4] to characterize search frictions and study the welfare impact of tariff price changes compared to a frictionless matching technology. While I characterize YT's search behavior exactly as in Buchholz (2022), I use it to analyze an additional choice of technology rather than changes in YT operations.

Frechette et al (2019) include a similar analysis where they assess the impact of the YT fleet being divided equally amongst two hail modes: street-hail and a central dispatcher with an Uber-like technology. They show that while the central dispatcher is better for those who are matched by it, the resulting market segmentation increases aggregate wait times and reduces demand and entry, especially when the market is denser (during the day). I differ from their approach in two ways. First, I allow YT to endogenously choose when to use the Uber app rather than assuming that a fraction of the drivers do; second, Frechette et al (2019) focuses more on aggregate entry/exit whereas my focus is on individual drivers' optimal search behavior.

Shapiro (2020) also focuses on Ubers' optimal search behavior to character-

[4]The author mentions that this model adapts elements from Lagos (2003)'s spatial search model.

ize density impacts for Ubers' welfare benefits. Building on the framework by Buchholz (2022), this paper estimates that the consumer surplus benefits could be 10 times more in the less dense areas. Employing a counterfactual where all Ubers are replaced by YT, the paper argues that this is due to lower wait times facilitated by Uber's better matching technology.[5] In more dense areas, street-hail wait times are comparable to Ubers' and the welfare gains disappear. I use Shapiro (2020)'s approach to approximate Uber's matching technology below. I find similar results for the boroughs, but I also find a substantial shift to the app for YT even in the dense areas. Unlike Shapiro (2020), my focus is on how drivers utilize the different technology along with how consumers react to it.

Ghili and Kumar (2020) add an explicit analysis of market 'thickness' (i.e. platform size) on drivers' spatial choices to estimate a threshold of platform size below which the supply of cabs is highly concentrated in high density zones while low density zones are under served.[6] As mentioned above, by allowing drivers' to choose matching technology (and hence platform size) endogenously, the reverse question can also be tackled here, i.e., which platforms do drivers use when they're in a dense region.

Although the agreement between Uber and YT isn't a merger, it could also potentially have market power and platform competition related implications if a majority of YT drivers switch to Uber. For example, by focusing on the shipping industry, Brancacchio et al (2020a) assert that the market power of a centralized platform erodes part of the gains from less search frictions. Similarly, Rosaia (2020) finds that a merger between Uber and Lyft would substantially increase overall welfare by 'thickening' the market, but the gains would be entirely captured by the monopolist. Another relevant element of this, and of Brancacchio et al (2020b), is to model price setting mechanisms within a search model by using the valuations at origin and destination. As I mentioned above, I don't model Uber prices in this paper but I could adopt similar strategies in the future to do so.

Since the focus is drivers' hail-mode choice, the paper further contributes to the theory of multi-homing in the transport literature. Bryan and Gans (2019) and Loginova et al (2020) relate social optimality of platform market power to multi-homing behavior. Tadepalli and Gupta (2020) find that multi-homing could lead to an efficient platform subsidizing drivers' time on the inefficient one. Yu et al (2021) use Machine Learning to identify driver characteristics that impact multi-homing behavior. To the best of my knowledge, this literature does use structural IO techniques to analyze multi-homing. The closest study is He et al (2018) who

[5]There are obviously other innovations that make Uber relevant beyond just the matching technology. For example, Cramer and Kruger (2016) relates Uber's generally higher utilization rates in part to their flexible supply model, Castillo (2020) finds increased consumer surplus due to surge pricing compared to a high average uniform price, and Bimpikis et al (2019) show that under 'unbalanced' demand (e.g. more rides to a place than from it), pricing rides differently based on origin can be optimal.

[6]In the context of the 'Wild Goose Chase' described in Castillo et al (2017), this could perhaps be understood as the drivers endogenously setting a 'small maximum dispatch radius'.

build a network market equilibrium model and empirically test optimal pricing and penalty strategies by assuming a set of model parameter values.

By employing the computation strategy of Buchholz (2022) and Rosaia (2020), this paper also relates to the literature on estimating dynamic models where agents play strategies without full information on payoff-relevant state variables as in Weintraub et al (2008) and Fershtman and Pakes (2012). The local matching function approach used for YT also relates to Brancacchio et al (2019).

The rest of the paper is organized as follows: in section I, I provide more details about the regulatory environment of the taxicab market in NYC, discuss the data, and provide descriptive statistics. Section II describes the demand and supply models, followed by the estimation strategy in section III. I present the estimation results in section IV followed by the counterfactual in V. Section VI concludes. The tables and figures are in the Appendix, A.

I. Market Description

A. Regulatory Environment

The medallion system keeps YT supply limited (currently to 13,587 active ones), while their fares also fixed (except for some overnight and rush hour surcharges). YT drivers need to have a special TLC permit to drive, but they're allowed to pickup passengers from across NYC. On the other hand, they are prohibited from declining rides once they've been hailed by a passenger. YT drivers work in either morning or evening shifts, and must drive a minimum number of hours in a year specified in the TLC Rules.

While there are no restrictions on their supply, Uber (and Lyft) drivers also need TLC permits to operate and are linked to HVFHS bases. Their fares are not fixed and while they are allowed to pickup passengers from anywhere in the city, they are not allowed to street-hail. Although Uber drivers can decline pickup requests on the app, Uber's Cancellation Policy page mentions that they could restrict access to the platform for drivers who breach the cancellation policy. There are no constraints on their driving duration and there are no shifts.

An important question is whether yellow taxi drivers could use the Uber app even before this agreement. TLC rule $\S 59A - 09(b)$ prohibits one vehicle from being linked to multiple types of licences, while TLC rule $\S 80 - 02(1)$ restricts drivers from operating a vehicle not associated with their driving permit type. So it is safe to assume that wasn't happening before September 2022 which is when the feature finally rolled out[7].

Two other features of the market are ignored in the paper. First, I don't include green taxis below because they are excluded from Manhattan and have negligible presence in the borough zones I consider. Second is the existence of two e-hailing platforms for yellow taxi which Uber is partnering with to list YT on

[7]Source: Jackie Davalos (Sep 7, 2022), "Uber Begins Rolling Out Yellow Taxi Rides in New York City," Bloomberg

their platform: Arro and Curb. Although YT were available on these apps before the agreement, the proportion of rides booked through them remains small[8]. Since TLC doesn't identify which YT rides are booked through an app, I assume all of them happened via street-hail instead.

B. Data

I use the TLC trip record data from March, 2022, right before the announcement was made, for the analysis below. Table A1 mentions the available information for both yellow taxi and Uber, as well as the total number of rides for the month.[9]

While this data is quite popular in the literature, TLC no longer releases Taxi IDs which were used in previous studies to link one trip to another.[10] The pickup and dropoff locations are also aggregated at the TLC zone level rather than longitude and latitudes. These changes imply that yellow taxi wait times can no longer be estimated using the strategies adopted by Frechette et al (2019) and Shapiro (2020). The lack of taxi identifiers is a bigger issue as it precludes me from identifying entry/exit behavior. On the other hand, I observe actual wait times for Uber and Lyft trips as well as the exact payment to the drivers which wasn't available before 2019.

C. Descriptive Statistics

Figure A1 shows the total count of pickups for YT and Uber across zones aggregated at the month level (which are qualitatively no different across hours of the day).[11] It's quite clear that YT pickups are concentrated heavily in the dense regions of NYC (Manhattan and the airports), whereas Ubers are more spread out. Figure A2 plots an indicator function of drop-offs vs pickups, $dp_{ind} = \mathbf{1}(d > p) - \mathbf{1}(d < p)$, for both cab types. From this, we can infer that (unlike Uber), if YT drops a passenger off in one of the less dense regions, it's far more likely to drive vacant to one of the dense regions looking for the next passenger.

Figure A3 adds more context by plotting drop-offs vs pickups for both cab types in a scatter plot. It shows that while Uber pickups and drop-offs more or less align with each other, that's not true for YT. However, the level of this effect, which I term 'utilization inefficiency' from the perspective of a YT driver, is small. Figure A4 zooms into the yellow taxi scatter plot by conditioning on the number of pickups. It becomes clear that while most NYC taxi zones have very few YT pickups (<2,000 pickups in 75% zones), *all* of these zones also have substantially more drop-offs. Put another way, while it's not very common for YT to drop a

[8]Rides booked through Curb saw a spike after the pandemic, but it still only accounts for 10% of all daily rides. Source: Preetika Rana (March 24, 2022), "Uber Reaches Deal to List All New York City Taxis on Its App," *The Wall Street Journal*

[9]The number of YT rides in a month, 3.6 million, is precisely the minimum platform size that Ghili and Kumar (2020) estimated below which supply converges to high density zones.

[10]A Freedom of Information Law (FOIL) request to the TLC was declined on anonymity concerns.

[11]The average number of rides and market shares over time can be found at a blog by Todd W. Schneider, Taxi and Ridehailing Usage in New York City.

passenger in one of the boroughs, but if they do so, they almost surely *have* to drive vacant till they reach another dense zone.

Ghili and Kumar (2020) show that such a scenario (drop-offs > pickups consistently) can be explained by a situation where people have learnt not to look for the platform because it is too under-supplied. Shapiro (2020) also showed Uber's welfare gains in the boroughs were due to lower wait times which induced demand rather than substitution away from another platform. Both these imply that the 'utilization inefficiency' is driven by lack of demand for YT as opposed to some unobserved preference of YT for Manhattan. To test this hypothesis in a simple manner, I match Uber and yellow trips in the data, i.e., find the trips on both platforms which have the same origin-destination, pickup times within 15 seconds, drop-off times within 30 seconds, and trip distance within 0.1 mile of each other. Using around 3,300 such trips that I find in the data, I plot figure A5 to show that, in fact, Uber rides were almost always more expensive than a YT, conditional on everything else being almost the same.

This is relevant here because any such unobserved preference, if it existed, would have to be modelled inversely to the better match technology offered by Uber in the counterfactual. Moreover, this makes it intuitively clear that, if nothing else, having the app would help YT pick up passengers in a borough after a dropoff.

II. Model

This section describes the model I use to compute the dynamic spatial equilibrium. I use a nested logit specification for demand and follow the work of Buchholz (2022) and Shapiro (2020) for supply.

A. *Demand*

For any time $t \leq T$ where T is the final period being considered in a day, let θ_{ij}^t denote the number of passengers who wish to travel from location i to j. I assume their choice of travel mode can be grouped into two nests: cab (yellow taxi, Uber, or Lyft) and public transport (bus or subway). Similar to the strategy adopted by Berry et al (1996) and Rosaia (2020), I treat the goods in one nest (public transport) as the outside option, and use the features of the goods in the other nest (cab) for the estimation. I treat each i-j route at a time t as a market. With this setup, the utility of an individual choosing travel mode $h \in (y, u, l)$ where y is yellow taxi, u is Uber, and l is Lyft to go from i to j at t is given by,

$$U_{ij}^{h,t} = \alpha q_{ij}^{h,t} + \beta w_i^{h,t} + \gamma \mathbb{1}\{h \in (u,l)\} + \xi_{i,j}^{h,t} + \zeta_{i,j}^t + (1-\sigma)\epsilon_{i,j}^{h,t} \tag{1}$$

Here, $q_{ij}^{h,t}$ is price for the trip on mode h, $w_i^{h,t}$ is the expected wait time in location i for mode h, $\mathbb{1}\{h \in (u,l)\}$ is an indicator which equals 1 for modes which are hailed via an app (Uber and Lyft) and $\xi_{i,j}^{h,t}$ are the unobserved (to the researcher) characteristics for product h that may be correlated to price or wait times. $\sigma \in (0,1)$ is a measure of correlation in the unobserved utility among the cab types; if $\sigma = 0$, the model reduces to a standard logit. $\zeta^t_{i,j} + (1-\sigma)\epsilon^{h,t}_{i,j}$ is chosen to yield the nested logit market shares (McFadden (1978)).

Using the nested logit shares, I can define the number of passengers who prefer mode h at t to go from i to j as,

$$\lambda^{h,t}_{ij} = \theta^t_{ij} s^{h,t}_{ij} \tag{2}$$

B. Matching and Wait Times

Yellow Taxi

As in Buchholz (2022), I use a local matching function to define how YT match with their passengers. Assume that there are $v^{y,t}_i$ vacant YT in a region i at time t and $\lambda^{y,t}_i$ passengers looking for them. The matching function[12] relates these up to a region-specific efficiency parameter, α_r as,

$$m_r(\lambda^{y,t}_i, v^{y,t}_i) = v^{y,t}_i \left[1 - e^{-\frac{\lambda^{y,t}_i}{\alpha_r v^{y,t}_i}}\right] \tag{3}$$

This implies an associated probability that a taxi in location i at t finds a match,

$$p^{y,t}_i = \frac{m_r(\lambda^{y,t}_i, v^{y,t}_i)}{v^{y,t}_i} = 1 - e^{-\frac{\lambda^{y,t}_i}{\alpha_r v^{y,t}_i}} \tag{4}$$

Similarly, the probability of a passenger (denoted by c) in i at t finding a taxi is given by,

$$p^{c,y,t}_i = \frac{m_r(\lambda^{y,t}_i, v^{y,t}_i)}{\lambda^{y,t}_i} \tag{5}$$

With this, a survival function can be used to approximate expected wait times (in terms of number of periods) as in Shapiro (2020). The probability of matching in period x after a passenger arrives at t implies expected wait times as,

[12]This aggregation matching function is based on the urn-ball specification from Hall (1979)

$$S_i^{c,t}(x) = p_i^{c,y,x+t} \prod_{k=1}^{x-1} (1 - p_i^{c,y,k+t}) \implies w_i^{c,t} = \sum_{k=1}^{\bar{T}} k S_i^{c,t}(k) \tag{6}$$

where $x + t \leq T$ in the survival function and $t + \bar{T} \leq T$ for the wait time. Essentially, the wait time is the expected value of the survival function which considers the match probability in a period x after not matching before.

This specification captures the inherent search frictions in the market where taxis and passengers don't necessarily match even when they're looking for each other. The underlying assumptions here are that matches can only occur when both the cab and passenger are in the same location, they're randomly assigned to each other, and once a match is made, the drivers can't refuse the ride.

Uber

I follow the approach described by Shapiro (2020) for the matching process of Ubers. Rather than using an explicit algorithm which tries to replicates Uber's technology, it approximates the process by capturing some of its key features. For instance, there are two key differences between the matching technology of YT and Ubers - Ubers need not be in the same location as the passenger to get matched, and a passenger booking a ride through the Uber app would almost certainly find a match. The first of these is modeled below, and the second is assumed to be true with certainty, i.e., no passenger who looks for an Uber remains unmatched. This certainty in finding a match, or lack of search friction, is justified by accounting for the variation in wait times determined by the vacant Uber distribution. Explicitly, every passenger will match with an Uber, but the wait times would be high or low depending on where the cabs are located.

Consider $\lambda_k^{u,t}$ passengers in k who are looking for an Uber t. The share of passengers who match with an Uber in location i is given by a logit probability,

$$p_{i,k}^{c,u,t} = \frac{\exp(v_i^{u,t}/(1+\tau_{ik})\mathbb{1}(\tau_{ik} \leq \bar{\tau}))}{1 + \sum_{i'} \exp(v_{i'}^{u,t}/(1+\tau_{i'k})\mathbb{1}(\tau_{i'k} \leq \bar{\tau}))} \tag{7}$$

where $v_i^{u,t}$ are the vacant Ubers in i and τ_{ik} is the time it would take for an Uber to travel from i to k. This specification implies that closer regions with more vacant cabs are likely to be the source of a match. $\bar{\tau}$ limits the distance up to which Uber would allow a match between a passenger and a cab. From this, the probability of an Uber in i matching with a passenger in k is given by,

$$p_{i,k}^{u,t} = \frac{\lambda_k^{u,t}}{v_i^{u,t}} p_{i,k}^{c,u,t} \tag{8}$$

Note that $\sum_k p_{i,k}^{u,t} v_i^{u,t} = \lambda_k^{u,t}$, i.e., each passenger finds a match. However, this doesn't preclude an Uber from remaining unmatched. Assuming further that the Uber app displays wait time using the process above, the expected wait time for a passenger in k can be approximated by the expected time it takes for a potentially matched cab in i to reach k,

$$w_k^{c,u,t} = \sum_i \tau_{ik} p_{i,k}^{u,t} \tag{9}$$

C. *Supply*

At any given time, cab drivers' behavior depends on their own location, whether they are employed or vacant, and the state of the world, $\mathcal{S}^t$. The state of the world is defined by the location of both employed and vacant cabs.

Conditional on the state, the ex-ante value of a vacant YT driver a at location i in time t is given by,

$$\begin{aligned} V_i^{y,t}(\mathcal{S}^t) =& \mathbb{E}\Big[p_i^{y,t} \underbrace{\Big(\sum_j M_{ij}^{y,t} \cdot (\Pi_{ij}^{y,t} + V_j^{y,t+\tau_{ij}}(\mathcal{S}^{t+\tau_{ij}})) \Big)}_{\text{Exp. value of ride}} + \\ & (1-p_i^{y,t}) \underbrace{\mathbb{E}\Big(\max_{j \in A(i)} \big\{ \overbrace{V_j^{y,t+\tau_{ij}}(\mathcal{S}^{t+\tau_{ij}}) - c_{ij}^y}^{W_i^{y,t}(j)} + \epsilon_{j,a}^y \big] \Big)}_{\text{Exp. value from vacancy}} \Big] \end{aligned} \tag{10}$$

The first term describes the value that a vacant driver receives from matching with a passenger at location i and dropping the passenger off at j. $p_i^{y,t}$ is the probability of matching defined by (4). $M_{ij}^{y,t}$ denotes the transition probability of an employed cab, i.e., the probability that a passenger in i that gets matched with the vacant cab wants to go to j. $\Pi_{ij}^{y,t}$ is the trip profit, i.e., driver pay subtracted by the fuel cost of driving the passenger from i and j at time t. $V_j^{y,t+\tau_{ij}}$ is the continuation value in j at $t+\tau_{ij}$ where τ_{ij} captures the time taken to drive there.

The second term describes the choice variable for the driver, that is, which zone j they search for passengers in if they are not matched in period t. $A(i)$ are the possible search locations near i (including i) which the drivers can reach within the next period. c_{ij}^y captures the associated costs of travelling from i to j and the driver-specific shock makes up the ex-ante unobserved component of the expected profit from this choice.

For Uber, the ex-ante value of a vacant driver a at location i in time t is,

$$V_i^{u,t}(\mathcal{S}^t) = \mathbb{E}\Big[\sum_k p_{i,k}^{u,t} \underbrace{\big(\sum_j M_{kj}^{u,t} \cdot (\Pi_{kj}^{y,t} - c_{ik}^u + V_j^{u,t+\tau_{ik}+\tau_{kj}}(\mathcal{S}^{t+\tau_{ik}+\tau_{kj}}))\big)}_{\text{Exp. value of ride from } k} + \tag{11}$$

$$(1 - \sum_k p_{i,k}^{u,t}) \underbrace{\mathbb{E}\Big(\max_{j \in A(i)} \big\{ \overbrace{V_j^{u,t+\tau_{ij}}(\mathcal{S}^{t+\tau_{ij}}) - c_{ij}^u}^{W_i^{u,t}(j)} + \epsilon_{j,a}^u \big] \Big)}_{\text{Exp. value from vacancy}} \Big]$$

The first term captures the value that a driver in i receives by matching with a passenger in k (with $p_{i,k}^{u,t}$ from (8)) who wants to travel to j. Therefore, there is an additional driving cost from i to k, c_{ik}^u, on top of the cost included in the trip profit from k to j. The other terms in the equation are as described above.

The key question for drivers of both cab types is where to search in case they aren't matched. Since this process is similar for both, I drop the y and u superscripts to simplify notation. Now, the ex-ante choice-specific value function of choosing to search in j at $t + \tau_{ij}$ for driver a before observing $\epsilon_{j,a}$ can be written as,

$$Z_i^t(j_a, \mathcal{S}^t) = \mathbb{E}[W_i^t(j_a)] = \mathbb{E}\big[V_{ja}^{t+\tau_{ija}}(\mathcal{S}^{t+\tau_{ija}}) - c_{ija}\big] \tag{12}$$

Assuming that $\epsilon_{j,a}$ are distributed according to a Type-I Extreme Value distribution allows the choice of j to be written using the familiar logit shares. That is, the probability that a vacant driver a chooses $j \in A(i)$ to search in the following period is given by,

$$P_i^t[j_a|\mathcal{S}^t)] = \frac{\exp\big(Z_i^t(j)/\sigma_\epsilon\big)}{\sum_{k\in A(i)} \exp\big(Z_i^t(j)/\sigma_\epsilon\big)} \tag{13}$$

where σ_ϵ captures the weight that unobserved ex-ante shocks carry in defining the drivers' location choice as opposed to their expectation of matches and the associated value. This equation defines the policy function, σ_i^t, that drivers use between state transitions to maximize their utility.

D. Equilibrium

The equilibrium is given by a sequence of state vectors (locations of employed cabs, $e_i^{h,t}$, and vacant cabs, $v_i^{h,t}$), policy functions (probability of vacant cabs searching in zone j, $\sigma_j^{h,t}$), transition beliefs (for both vacant and employed cabs) and an initial state $\mathcal{S}^{h,0}$, such that,

- YT matches happen according to the matching function in equation 3, Uber matches happen according to the mechanism in equations 7 and 8. Employed cabs move according to the respective transition matrix, $M_{ij}^{h,t}$, in both cases.
- Vacant cabs of both types in each location follow the policy function $\sigma_i^{h,t}$ to search in the following period given their expectation of state transition kernels,
- State transition kernels are defined by both employed and vacant cab movements,
- Drivers' beliefs are rational so the state transition beliefs of vacant cab drivers are self-fulfilling.

For more details on the equilibrium definition and a proof of its convergence, refer to Buchholz (2022) and Shapiro (2020).

III. Estimation Strategy

This section begins with a brief description of the subset of NYC zones that I use to estimate the model, followed by the strategy adopted to compute both demand and supply. I use weekday data (Monday to Thursday) from March 2022 for all the calculations below.

A. Subset of Zones

For both demand and supply, I use data from 14 neighboring TLC zones, 8 in Manhattan and 6 in the boroughs, where the Manhattan zones are connected to the borough zones by the Queens Midtown Tunnel. Essentially, I subset the entire month's data to consider rides which are taken only between these zones on weekdays. That leaves 136,990 YT trips and 236,237 Uber trips. These are about 3.8% and 1.8% of the total trips in the month.

I focus on the subset as a simplification, both from the perspective of computation and interpretation of results. Figure A6 shows the pickup counts in the subset zones. Figure A4 shows the distribution of drop-offs - pickups in the whole city and compares it with the distribution in the subset. These graphs show that the qualitative patterns of YT and Uber behavior in these zones are similar to the rest of the city. Specifically, YT are highly concentrated in the Manhattan zones and Uber has a more even distribution across both types of zones. Since the primary impact I'm trying to analyze is how the additional hail mode choice interacts with their location choice (which is dependent on density) for a YT driver, even though these are only a fraction of the total trips, the estimates from this subset would act as an approximation of the overall impact. In fact, as we consider boroughs farther from Manhattan, the impact of density on hail mode choice is likely to be higher as YT have even lower pickups in those zones.

B. Demand

Given the nested logit setup in (1), I normalize the mean utility from the outside goods to 0 and use the Berry inversion (Berry (1994)) to estimate the following equation using Linear IV GMM,

$$ln(s_{ij}^{h,t}) - ln(s_0^{h,t}) = \alpha q_{ij}^{h,t} + \beta w_i^{h,t} + \gamma \mathbb{1}\{h \in (u,l)\} + \sigma ln(s_{ij}^{h|g,t}) + \xi_{i,j}^{h,t} \tag{14}$$

where $s_{ij}^{h,t}$ is the market share of mode h and $s_{ij}^{h|g,t}$ is the share of mode h within the cab nest. I define a market as a 30-min interval for a route that can be between the combined Manhattan zones and the combined borough zones. That is, t is a 30-min interval and (i,j) can be one of $(M,M),(M,B),(B,M),(B,B)$ where M Manhattan and B denotes the borough. Given the discussion in I.C, I allow the coefficients to vary by pickup area, i.e., the estimates for $(M,M),(M,B)$ differ from $(B,M),(B,B)$.

Shares

For the outside goods, I use both the MTA Travel Survey data from 2019 and the MTA daily ridership data for March 2022. The daily ridership data provides the total number of riders in NYC buses and subways on each day of the month. However, it doesn't provide details of the routes and times of those rides. To approximate these, I assume that if N_d^b and N_d^s are the total number of riders in a bus and subway on day d, the number of those riders going from i to j at t can be approximated by $N_d^b\pi_{ijt}^b$ and $N_d^s\pi_{ijt}^s$ where the shares don't change by day. While this is a strong assumption, it is not uncommon in the literature (see Shapiro (2020) and Rosaia (2020)) and can be somewhat rationalized by assuming low variation in morning weekday commute patterns.

I compute these fixed shares, π_{ijt}^b and π_{ijt}^s, using the latest MTA Travel Survey (March to June 2019) which provides peoples' travel mode choice details up to the census tract level at different times of the day. I match the census tracts with TLC zones and construct π_{ijt} by calculating the share of trips at each market level against the total observed trips for that mode (bus or subway) in the survey. This is the step which forces me to define markets up to the region (M or B) instead of the TLC zone level because I encounter 0 outside shares otherwise. This is also the reason that I club buses and subway together as the outside nest. This could be an issue if they were very different from each other, but they have the same price in NYC ($2.75) and the aggregation up to the regions at least reduces, if not mitigates, localized impacts of different routes and timings.

Using $N_d^s\pi_{ijt}^s$ and $N_d^b\pi_{ijt}^b$ as the total number of passengers choosing the outside goods, I construct the cab shares by aggregating the total number of pickups in the TLC ride data at each market level.

YT wait times

While wait times for Uber and Lyft are observed in the data, that is not the case for YT. Previous studies like Frechette et al (2019) and Shapiro (2020) have estimated YT wait times by using algorithms which relate the the dropoff patterns of taxis, the street layout of NYC, and the average speed on those streets. While I don't have longitude and latitude data of taxi dropoffs to even attempt such an algorithm, luckily I was provided with a snapshot of the average wait times estimated by the author of Shapiro (2020) for 9 AM in NYC in Sep-Dec 2015. I simply regress these wait times on the relevant variables from their original algorithms using data from Sep-Dec 2015 to approximate a relation between the two and get,

$$w^t = \underset{(8.162)}{367.71} - \underset{(34.171)}{55.395}\mathbb{1}\{M\} - \underset{(0.165)}{0.6687}d^{t-30} - \underset{(614.389)}{2215.4}\, s^{t-30} + \underset{(12.458)}{24.552}d^{t-30}s^{t-30} \tag{15}$$

where w^t are the Shapiro (2020) wait times in seconds, $\mathbb{1}\{M\}$ is an indicator for downtown Manhattan zones, d^{t-30} are dropoffs and s^{t-30} is speed (miles/second) in the previous 30-minutes within each zone ($R^2 = 0.4132$). The speed is calculated using the distance travelled and time taken for the completed trips within the zone. The results are intuitive since more drop-offs (more available cabs) and a higher speed (faster to reach a passenger) are negatively related with wait times. Even though the coefficient on the Manhattan zones isn't statistically significant, it follows the intuitive idea from Shapiro (2020) that YT wait times in downtown Manhattan are lower than elsewhere. Therefore, I use the March 2022 data and the point estimates above to approximate wait times for the demand estimation. Figure A8 compares these with the Uber and Lyft wait times across regions. While the borough wait times are generally higher than Manhattan, there is clearly less variation in them than the observed values for Uber and Lyft due to the way they're constructed. While this isn't ideal, I hope to receive more details on the estimates from Shapiro (2020) in the future to better approximate the wait times.

Instruments

For Uber or Lyft price, I use the price in the same region from the previous period as an instrument. This is driven by the description of the Uber surge pricing mechanism in Castillo (2020). He shows that surge multipliers are smoothed out over space and time (with high auto-correlation in the first 10 minutes and settling around 0.15 up to 30 min). So, prices from the previous period would be correlated with current prices without depending on current demand thereby satisfying the exclusion restriction.

For wait times, I use the speed of the cabs calculated from completed trips within the area in the same period. Wait times depend on both the number of vacant cabs and how quickly they can get to a passenger. This instrument captures the second part of this mechanism in two ways. First, it captures the aggregate driving conditions in the zone which could vary exogenously over time due to weather, road closures, and other vehicles on the roads. Second, even assuming homogeneous driver quality, the speeds vary across cab types because of the differences in their navigation technology which is independent of demand conditions. While YT drivers may use Google Maps or navigate by experience, both Uber and Lyft have their own navigation systems and allow drivers to change them.[13] I don't use previous dropoffs in the zone because even though they are relevant for YT, the dropoffs for Uber and Lyft are less relevant since they don't need to be in the same region for a match.

For within group share of cab type h, I use the average prices of the other two cab types in the same period and same zone. If the prices of the other cabs are high, it is likely that cab type h would have a higher within group share. While the average prices of h_{-1} would be correlated with their own demand unobservables, they should be independent of the unobservable demand characteristics of h.

C. Supply

Since I closely follow Buchholz (2022) and Shapiro (2020) for the supply algorithm, I focus only on the salient features of the model and any changes I make specifically for the computation in this paper.

Assumption 1: Total number of yellow cabs, $\bar{y}$, and Ubers, $\bar{u}$, is fixed. There is no entry or exit.

This is driven by the ride data I use to compute the equilibrium. While the model could be extended to include entry (using an outside option while calculating the policy functions) and exit (using a cost disutility that increases with time, e.g. due to fatigue, in the value functions) as in Frechette et al (2019) and Rosaia (2020), these wouldn't be identified without taxi (or driver) identifiers which provide information on entry/exit patterns. The only information I have is from the TLC monthly reports which state that an average YT driver was active for 8.1 hours a day and an Uber driver for 6.6 hours in March 2022. I compute the equilibrium using data from 6 AM to 4 PM so it is quite likely Uber entry and exit would play a role. Assuming that away is one of the key limitations of this paper.

Assumption 2: Given $\bar{y}$, $\bar{u}$, observed matches, and average profits per ride, σ_ϵ is set to the minimum value that allows the search model to converge.

As mentioned above, σ_ϵ determines how much drivers' choice (doesn't) depend on their expectation of future values. A higher σ_ϵ would cause the choice probabilities in the policy function to tend towards a uniform distribution, which would increase drivers' spatial distribution. To accurately identify σ_ϵ, Buchholz (2022) and Shapiro (2020) use information on vacant cab behavior (e.g. utilization rates or average vacancy time between trips) and match the moments generated by their model to the data. Since I don't observe this information, I'm unable to follow a similar approach.

[13]Source: Gridwise (2021), "The best navigation apps for gig drivers"; Uber and Lyft's help pages for drivers.

Instead, I approximate a lower bound for σ_ϵ that relies on its relation with total cabs and expected profits. For a given $\bar{h}$, as σ_ϵ reduces, it forces the cabs to be highly concentrated in (near) high valuation zones. However, this skews the match probability towards low valuation zone, which causes the cabs to move there in the next iteration. So the algorithm keeps bouncing between states. By setting σ_ϵ to the minimum value that allows convergence, I essentially assume that the search behavior is driven by at least some level of ex-ante uncertainty instead of just expected valuation.

Assumption 3: Drivers have full knowledge of the initial state vector $\mathcal{S}^0$ and all model parameters including those from the demand side.

This assumption allows employment of an Oblivious Equilibrium concept (suggested by Weintraub et al (2008) and used by both Buchholz (2022) and Shapiro (2020)) where agents play strategies based on their expectations over industry averages rather than full knowledge of payoff-relevant state variables.

As in Buchholz (2022), for weekday data from 6 AM to 4 PM, I divide time into 5-minute intervals (on the lower end of the trip duration distribution) and estimate $\tilde{m}_i^t$ as a sixth-order polynomial of time to account for those periods with no pickups. Cost per trip are calculated by using mileage of the cab fleets and weekly fuel prices[14] along with the observed trip distance. I take monthly averages of profits and costs for trips from i to j in each 5-minute interval t. M_{ij}^t is calculated by using the average share of observed trips between each origin-destination pair in each time interval. I also set $A(i)$ as the zones that can be reached within two time periods based on the average trip duration between each origin-destination pair.

All this is done for both YT and Ubers separately. For Uber, I define $\bar{\tau} = 30$ minutes since 99.99% of the trips have lower wait times than that. With these values, I compute the algorithm presented in Online Appendix A.8 of Buchholz (2022) separately for both the cab types using their respective data, matching technology, and value function definitions. Since I don't make any changes in the algorithm, I don't mention the steps here.

This algorithm inherently models two countering forces of drivers' preferences, specifically, searching in regions with a) higher profits and continuation values, and b) higher match probability. If a driver goes to a an (a) type region, the policy function accounts for the lower match probability there in the next iteration. That makes the region less attractive and the policy function shifts to accommodate that. The balance of these two forces determines where the vacant cabs distribute and in what proportions.

Along with the values read from data described above, the algorithm requires the number of cabs and σ_ϵ to be specified. For Uber, I perform a grid search for $\bar{u}$ by re-computing the algorithm to match the mean wait times implied by the model (from 9) and observed in the data. This gives $\bar{u} = 499$. Since the matching

[14]Mileage from https://www.fueleconomy.gov/feg/findacar.shtml; fleet information from TLC (2020), Factbook; fuel prices from Weekly Average Motor Gasoline Prices in the NYC Metropolitan Area.

process assumes each passenger finds a cab, wait times can be directly related (inversely) to the number of cabs. I then perform another grid search for σ_ϵ to find the minimum value required for convergence and find $\sigma_\epsilon^u = 5.8$.

The same strategy can't be applied for YT since the wait times depend on the matching efficiency and potential passengers, both of which are unobserved. Instead, I approximate $\bar{y} = 364$ by dividing the average number of YT trips per day in my subset (7,203) with the average trips per vehicle per day (19.74) calculated from the TLC monthly report for March 2022. Performing a similar grid search gives $\sigma_\epsilon^y = 4.4$.

IV. Empirical Results

A. Demand

Table A2 provides the demand estimates for rides that originate in Manhattan and boroughs for 3 specifications. Specification (I) sets $\sigma = 0$, so the model becomes normal logit. Compared to specification (III) described in section III.B, (I) has lower standard errors for the borough estimates. Together with a statistically insignificant $\hat{\sigma}_B$ in (III) and a high estimate of $\hat{\gamma}_B$ in (I), I interpreted this as an indication that nesting all cabs together in the boroughs may not be appropriate. So in specification (II), I move YT from the "inside" nest to the outside options (similar to Rosaia (2020)) and remove the app indicator. $\hat{\alpha}_M$ is close to the Rosaia (2020) estimates for short trips, however $\hat{\beta}_M$ and $\hat{\sigma}_M$ are closer to those for medium distance trips. Part of the explanation for this difference could be that I have only used a subset of zones here. However, the borough estimates are once again imprecise.

Since Shapiro (2020) described close substitution patterns between Uber and YT in Manhattan, I use (III) which includes YT within the nest. $\hat{\sigma}_M$ is, in fact, close to 1 in Manhattan indicating very high correlation between the three cab types. $\hat{\alpha}_M$ also almost matches that of Shapiro (2020) for 150K-200K income range, but $\hat{\beta}_M$ is lower. Comparing Manhattan estimates and borough estimates from (III), I see close values for price but higher point estimates for wait time and app indicator in the boroughs. A higher $|\hat{\beta}_B|$ is likely explained by the wait time distribution in A8 since YT's wait times are towards the high end of Uber and Lyft's distribution while YT shares are significantly lower. The smaller shares for YT xould also explain the higher $\hat{\gamma}_B$.

While using the approximate YT wait times is definitely far from ideal, the point estimates seem to be in line with intuition and at least the Manhattan estimates are close to previous studies and are precisely estimated.

B. Supply

In this section, I show the vacant cab distributions from two representative zones in Manhattan and the boroughs for both cab types.

Figure A9 shows the vacant cab distribution for them in the same Manhattan zone. Figure A10 shows the vacant cab distribution in the boroughs. From these graphs, it's quite clear that the vacant cab patterns for Uber don't vary much across regions, whereas YT are way more likely to be in Manhattan than the boroughs.

A thing to note here is that the algorithm requires an assumption of the initial state. While that is not an issue for later periods in the day, none of the steps in the algorithm deal with updating the first time period. Due to this, the results for the first few periods (6-6.30AM) are driven more by the adjustments from the initial state rather than drivers' search behavior. I have also not estimated standard errors here. They could be estimated by re-sampling the data but I leave that for future work.

V. Counterfactual: Multi-homing

When YT are allowed to use the Uber app, a vacant YT driver would consider an additional choice variable in addition to the location for searching.[15] This could be defined as,

$$
\begin{aligned}
V_i^{y,t}(\mathcal{S}^t) =& \mathbb{E}\Big[p_i^{y,t} \underbrace{\big(\sum_j M_{ij}^{y,t} \cdot (\Pi_{ij}^{y,t} + V_j^{y,t+\tau_{ij}}(\mathcal{S}^{t+\tau_{ij}}))\big)}_{\text{Exp. value of ride}} + \\
& (1-p_i^{y,t}) \underbrace{\mathbb{E}\Big(\max_{j \in A(i), h \in (y,u)} \big\{ \overbrace{V_j^{h,t+\tau_{ij}}(\mathcal{S}^{t+\tau_{ij}}) - c_{ij}^{y}}^{W_i^{y,t}(j)} + \epsilon_{j,a}^{y} \big] \Big)}_{\text{Exp. value from vacancy}} \Big]
\end{aligned} \tag{16}
$$

This is different from (10) in the second term where the additional choice appears as $h \in (y, u)$ and $j \in A(i)$ is chosen using a policy function that accounts for continuation values of both YT and Uber. By not changing the first term, I'm implicitly assuming that vacant YT drivers only choose hail mode (simultaneously with search location) if they're not matched with a street-hail passenger at their location i in period t. Once they switch mode, they can either find an Uber pickup at j in period $t + \tau_{ij}$ or make both location and hail mode choices again.

The welfare impact of this additional choice would depend on not only the resulting vacant YT distribution, but also the vacant Uber distribution, Uber prices, and consequently new demand conditions. Since I haven't modelled Uber's price setting mechanism in this paper, I abstract away from those changes and

[15]One could argue that they don't have to choose between one or the other. That is, that they could multi-home on both hail modes simultaneously. However, as shown in Rosaia (2020) for Uber and Lyft drivers, while almost 50% drivers use both platforms, they don't do so simultaneously.

make a simplifying assumption that Uber prices and driver pay remains constant for each i and t. While this can be relaxed at a later stage using a different approach (like Rosaia (2020) or Brancacchio et al (2020b)), for now it remains another key limitation of this paper. I do, however, account for the response of vacant Uber drivers and allow demand to change at least partly due to changes in wait times. Algorithm 1 describes and explains the specific steps I take to compute the counterfactual while accounting for these changes.

Essentially, it has an inner 'Uber reaction' loop where the YT switching to the app are found by extending YT's policy function to include the continuation value for Ubers in each (i, t). Then, the YT who switch to Uber reduce existing Ubers' probability of matching so their distribution is re-optimized. Once this converges, the model implied wait times are used to change demand in the outer 'demand loop' and the whole process repeats till convergence.

I present results for the baseline model, a counterfactual with only the inner 'Uber reaction' loop (CF1), and finally for the full counterfactual including the 'demand change' loop (CF2). Table A3 shows the average number of vacant cabs across the day using different hail modes in different regions. While it appears that Uber drivers' aggregate search behavior doesn't change across regions, figure A12 shows that there are some small variations in individual zones of both regions.

The total YT from Manhattan reduces, and the remaining YT are almost evenly split on their hail mode choice. Around 15% of the cabs who were initially searching in Manhattan shift to the boroughs and use the app, and around 20% of the cabs who were searching in the boroughs adopt the app as well. Figure A11 plot these aggregates in each region over time for both YT hail modes.

This additional YT supply reduces Uber wait times in both the regions. On average across the day, they reduce by 42 seconds in Manhattan and 67 seconds in the boroughs. On the other hand, the wait times for YT (or more clearly, wait times for passengers seeking a street hail YT) increase by about 69 seconds in Manhattan, and only by 3 seconds in the boroughs.

Table A4 presents the welfare impacts of both CF1 and CF2 for each region along 3 measures. The first are changes in consumer surplus (CS) due to wait time shifts calculated using the log-sum approach for nested logit. They are presented without using the income disutility term, so it's appropriate to consider the ratio of CS rather than the numbers. The second is realized profits, i.e., the product of matches and average profits from a ride, i.e., $m_i^{h,t} \sum_{ij}^{h,t} M_{ij}^{h,t} \Pi_{ij}^{h,t}$. The third is the value function, $V_i^{h,t}(\mathcal{S}^{\sqcup})$. The values presented are totals across all periods of the day. For YT, profits are shown only for street hail matches, whereas their 'value' includes the full continuation value of hail mode choice.

In Manhattan, the CS remains virtually the same. The impact of lower Uber wait times seems to offset that of the higher street hail YT wait times. This is reflected in the observed profit variation for YT and Uber as well. Since I don't allow price to change, the only way profits are affected is through the number of matches. So, as YT start using the app, the higher wait times for street hailing causes substitution towards Uber, which in turn explains the substantially high uptake of the app by YT in Manhattan. On the other hand, comparing the value between Uber and YT reveals that the indirect utility from searching reduces substantially for Uber drivers compared to YT drivers. That is because of their lower probability of matching as YT enters the supply, whereas YT can pick up both street hail or app passengers while Ubers can't. In fact, since almost half the drivers who use the Uber app in Manhattan are YT, the profit increase for Uber is at least partly driven by the additional matches accruing to YT drivers rather than the original Uber drivers.

In the boroughs, CS increases by about 3 times. This is driven by lower Uber wait times as more cabs appear on the app in the region (both more uptake within the zones and additional YT who enter from the boroughs), whereas wait times for street hail YT barely increase. Part of the reason is that the demand in the borough was initially so low (for the reason discussed in I) that YT's leaving the street-hail market don't really have too much impact. Similar to Manhattan, the value for YT increases substantially while that of Uber decrases, whereas the profits for Uber (including erstwhile YT) increase substantially.

Comparing CF1 and CF2 shows us how changes in demand counteract the impact of the additional choice value for YT. Without changes in demand, the computation would over-estimate the impact since it wouldn't account for the lower matches accruing to street hail YT.

In aggregate, the shift increases CS (by reducing wait times) and increases profits for the Uber hail cabs by increasing demand. However, it is detrimental to the YT who continue to street hail.

VI. Conclusion

In this paper, I analyze the agreement between Uber and YT to add the YT on the Uber platform in NYC. For a subset of zones in NYC, I estimate a nested logit demand and compute a dynamic spatial supply equilibrium to run a counterfactual where YT can choose their hail mode and location. The counterfactual includes the impact of Ubers reacting to YT and change in demand due to wait times.

I find that a substantial number of YT choose the app in both Manhattan and the boroughs. This reduces wait times for Uber passengers and almost doubles overall consumer welfare, while increasing overall profits as well. The regional variations of these results are important and imply that while YT using the app would be better off, those who keep using street-hail would be worse off. However, these results suffer from two key limitations - lack of entry/exit and price response. I hope to work on this in the future and add these using perhaps a different dataset.

REFERENCES

Berry, S. (1994), "Estimating Discrete-Choice Models of Product Differentiation," *The RAND Journal of Economics*, 25 (2) 242-262.

Berry, S., Carnall, M., Spiller, P. (1996), "Airline Hubs: Costs, Markups and the Implications of Customer Heterogeneity," *NBER Working Paper*, 5561.

Bimpikis, K., Candogan, O. and and Saban, D. (2019). "Spatial Pricing in Ride-Sharing Networks." *Operations Research*, 67(3):744-769.

Brancaccio, G., Kalouptsidi, M. and Papageorgiou, T. (2019), "A Guide to Estimating Matching Functions in Spatial Models". *International Journal of Industrial Organization*, 70.

Brancaccio, G., Kalouptsidi, M. and Papageorgiou, T. (2020), "Search Frictions and Efficiency in Decentralized Transport Markets." *NBER Working Paper*, 27300.

Brancaccio, G., Kalouptsidi, M. and Papageorgiou, T. (2020), "Geography, Transportation, and Endogenous Trade Costs." *Econometrica*, 88(2), 657-691.

Bryan, K.A. and Gans, J.S. (2019), "A theory of multihoming in rideshare competition." *J Econ Manage Strat*, 28: 89 - 96.

Buchholz, N. (2022), "Spatial Equilibrium, Search Frictions, and Dynamic Efficiency in the Taxi Industry." *Review of Economic Studies*, 89(2), 556-591.

Castillo, J. C. (2020), "Who Benefits from Surge Pricing?" (Working Paper).

Castillo, J. C., Knoepfle, D. and Weyl, G. (2017), "Surge Pricing Solves the Wild Goose Chase." *Proceedings of the 2017 ACM Conference on Economics and Computation*, Pages 241–242.

Cohen, P., Hahn, R., Hall, J., Levitt, S. and Metcalfe, R. (2016), "Using Big Data to Estimate Consumer Surplus: The Case of Uber," *NBER Working Paper*, 22627.

Cramer, J. and Krueger, A. (2016), "Disruptive Change in the Taxi Business: The Case of Uber." *American Economic Review*, 106(5): 177-182.

Fershtman, C. and Pakes, A. (2012), "Dynamic Games with Asymmetric Information: A Framework for Empirical Work." *The Quarterly Journal of Economics*, 127, 1611–1661.

Frechette, G. R., Lizzeri, A. and Salz, T. (2019), "Frictions in a Competitive, Regulated Market: Evidence from Taxis." *American Economic Review*, 109, 2954–2992.

Ghili, S. and Kumar, V. (2020), "Spatial Distribution of Supply and the Role of Market Thickness: Theory and Evidence from Ride Sharing." *Cowles Foundation Discussion Papers*, 2562.

Hall, R. E. (1979), "A Theory of the Natural Unemployment Rate and the Duration of Employment." *Journal of Monetary Economics*, 5, 153–169.

He, F., Wang, X., Lin, X. and Tang, X. (2018), "Pricing and penalty/compensation strategies of a taxi-hailing platform." *Transportation Research Part C: Emerging Technologies*, 86, 263-279.

Lagos, R. (2003), "An Analysis of the Market for Taxicab Rides in New York City." *International Economic Review*, 44, 423–434.

Loginova, O., Wang, X. and Liu, Q. (2022), "The impact of multi-homing in a ride-sharing market." *The Annals of Regional Science*, 69. 1-16.

McFadden, D. (1978), "Modelling the Choice of Residential Location." *Transportation Research Record*, 672, 72-77.

Rosaia, N. (2020), "Competing Platforms and Transport Equilibrium: Evidence from New York City." *mimeo, Harvard University.*

Shapiro, M (2020), "Density of Demand and the Consumer Benefit of Uber." *Research Collection of Economics*, 1-74.

Tadepalli, K. and Gupta, A. (2020), "Multihoming in Ridesharing Markets: Welfare and Investment." (Working Paper).

Weintraub, G. Y., Benkard, C. L. and Van Roy, B. (2008), "Markov Perfect Industry Dynamics with Many Firms." *Econometrica*, 76, 1375–1411.

Yu, J., Mo, D., Xie, N., Hu, S. and Chen, X. (2021), "Exploring multi-homing behavior of ride-sourcing drivers via real-world multiple platforms data." *Transportation Research Part F: Traffic Psychology and Behaviour*, 80, 61-78.

Appendix

Algorithm 1 Algorithm for Computing Counterfactual

Set counter $k_d = 0$
Input wait time coefficients, $\hat{\beta}_M$ and $\hat{\beta}_B$.
Input $\hat{\alpha}_r$
Input initial $\lambda^{u,t}_{i,k_d}$ and $\lambda^{y,t}_{i,k_d}$.
repeat
    Set counter $k = 0$
    Label each zone according to taxi type, e.g. 170^y and 170^u.
    Set initial value for yellow taxis who switch to Uber, i.e., $v^{yu,t}_{i,k} = 0$.
    Input $\lambda^{u,t}_{i,k_d}$ and $\lambda^{y,t}_{i,k_d}$
    repeat
        Add $v^{yu,t}_{i,k}$ to each zone labelled Uber, z^u, for all t.
        Use $v^{yu,t}_{i,k}$ in $p^{u,t}_{i,k',k}$ calc, compute $v^{u,t}_{i,k}$ and $V^{u,t}_{i,k}(\mathcal{S}^t)\ \forall\ (z^u, t)$ for $\bar{u}$.
        Keeping $V^{u,t}_{i,k}(\mathcal{S}^t)$ fixed for z^u, compute $v^{y,t}_{i,k}$ and $\forall\ (z^u, z^y, t)$ for $\bar{y}$.
        Update $v^{yu,t}_{i,k+1} \leftarrow v^{y,t}_{z^u,k}$.
        $k \leftarrow k + 1$
    until $|V^u_k - V^u_{k+1}| + |V^y_k - V^y_{k+1}| < \delta\ \forall\ (i, t)$
    Compute wait times from model.
    Compute $\lambda^{u,t}_{i,k_d+1}$ and $\lambda^{y,t}_{i,k_d+1}$ using $\hat{\beta}_M$ and $\hat{\beta}_B$ and new wait times.
    $k_d \leftarrow k_d + 1$
until $|\lambda^u_{k_d} - \lambda^u_{k_d+1}| + |\lambda^y_{k_d} - \lambda^y_{k_d+1}| < \delta\ \forall\ (i, t)$

$\triangleright$

In the inner 'Uber reaction' loop, I compute Uber's continuation values, $V^{u,t}_i(\mathcal{S}^t)$, using the search algorithm. Keeping them fixed, I run the search algorithm for YT but include both $V^{y,t}_i(\mathcal{S}^t)$ and $V^{u,t}_i(\mathcal{S}^t)\ \forall\ (i, t)$ in the calculation of the policy functions $\sigma^{y,t}_i$ by modifying equation (13) to include the extra terms. This creates a set of $v^{yu,t}_i$ which are in i at t but due to $V^{u,t}_i(\mathcal{S}^t)$, not $V^{y,t}_i(\mathcal{S}^t)$. I then re-compute the vacant Uber distribution and value functions with the additional $v^{yu,t}_i$ included in the match probabilities. This repeats till convergence.

I use the resulting vacant cab distributions to calculate wait times (discussed below) from (6) and (9) in iteration k_d of the outer 'demand change' loop. With these wait times, I use my demand estimates to calculate the change in the nested logit shares of both Uber and YT. I scale equation (2) using the new shares (assuming a fixed θ^t_{ij}) and run the 'Uber reaction' loop for the new demand distribution.

Uber wait times can be calculated directly using the probability of matching computed in the 'Uber reaction' loop. To calculate YT wait times, I need $\hat{\alpha}_r$ so I can use the new vacant cab distribution to calculate (5) which appears in (6). As seen from A8, the average YT wait times are < 5 min in Manhattan, and usually ≥ 5 min in the boroughs. I make the simplifying assumption that average YT wait times are distributed uniformly, which allows me to write $m_M(\lambda^{y,t}_M, v^{y,t}_M) = \lambda^{y,t}_M$ and $m_B(\lambda^{y,t}_B, v^{y,t}_B) = 0.5\lambda^{y,t}_B$, i.e., in a period, only half the YT passengers get matched in the boroughs, while all of them get matched in Manhattan. Buchholz (2022)'s estimates of search frictions found that unmet demand is only about 8% of the unmatched vacant cabs in a period, so this assumption isn't entirely without basis. Doing this on the vacant cab distribution from section IV.B allows me to invert α_r from equation (3) which remain fixed throughout the computation.

Table A1—: TLC Ride Data

Range	March, 2022	
Taxi Zones	263	
	Yellow Taxi	**Uber**
Number of Rides	3.6 million	13.1 million
Taxi ID	No	No
Ride request time	No	Yes
Payment	Fare amount	Fare amount & driver pay
Pickup/Drop-off Location	By zone	By zone
Pooling	Passenger count	Shared ride request flag

Note: The data on pickup/drop-off time, trip distance, tips, additional fee/taxes/charges are also available for both.
Source: TLC Trip Record Data.

Table A2—: Demand Estimates

	(I) Logit	(II) Nested Logit	(III) Nested Logit
Include YT	✓	×	✓
α_M	-0.0081*	-0.0257***	-0.0279***
	(0.0036)	(0.0039)	(0.0025)
β_M	-0.0096***	-0.0105***	-0.0030***
	(0.0003)	(0.0005)	(0.0005)
γ_M	0.1825***	-	0.1844***
	(0.0476)		(0.0333)
σ_M	-	0.6183***	0.9880***
		(0.1112)	(0.0804)
α_B	-0.0232**	-0.0963***	-0.0252**
	(0.0083)	(0.0068)	(0.0096)
β_B	-0.0158***	-0.0014	-0.0140*
	(0.0004)	(0.0014)	(0.0060)
γ_B	2.0454***	-	1.7582
	(0.1001)		(0.9716)
σ_B	-	0.4237	0.1288
		(0.3753)	(0.4236)

Note: Specification (I) assumes $\sigma = 0$, i.e., the model is logit; (II) considers yellow taxi as an outside option and removes the app indicator; (III) is described in section III.B. M denotes Manhattan pickup areas, B denotes borough pickup areas.

Table A3—: Counterfactual Vacant Cabs

	Baseline		CF1			CF2		
	YT	Uber	YT		Uber	YT		Uber
Hail Mode	Street	-	Street	App	-	Street	App	-
Manhattan	191	194	76	90	195	78	84	195
Boroughs	84	142	66	43	141	64	48	141
Total	275	336	142	133	336	142	132	336

Note: These are average vacant cabs searching in a region across the whole day. CF1 includes only Uber's reaction while CF2 includes the change in demand.

Table A4—: Counterfactual Welfare

			Baseline	CF1 (Uber Reacts)		CF2 (Demand Changes)	
Manhattan	CS		143.23	-	-	143.06	0.99
	Profit	YT (SH)	49,078.1	-	-	34,848.8	-28%
		Uber	71,551.0	-	-	85,011.1	18%
	Value	YT	101,661.8	269,005.2	62%	141,130.8	38%
		Uber	147,374.1	104,541.3	-40%	106,665.8	-27%
Boroughs	CS		113.41	-	-	344.49	3.03
	Profit	YT (SH)	1,000.9	-	-	1,002.4	0.15%
		Uber	51,355.1	-	-	67,183.5	30%
	Value	YT	71,059.2	191,767.7	169%	116,340.6	63%
		Uber	109,969.1	78,148.3	-28%	79,657.5	-27%
Total	CS		256.64	-	-	487.55	1.89
	Profit		172,985.1	-	-	188,045.8	8%
	Value		430,061.2	643,462.5	49%	443,794.7	3%

Note: Consumer Surplus (CS) are calculated as log-sum values; they have not been divided by the marginal utility of income parameter. Profits are totals over all (i,t) of average profit in (i,t) multiplied by the number of matches. For YT, these are values for only the street hail (SH) mode matches. Value are totals of value function for all (i,t) across the day. For YT, these include the continuation values of location as well as hail mode choice. All percentages are calculated with respect to the respective baseline.

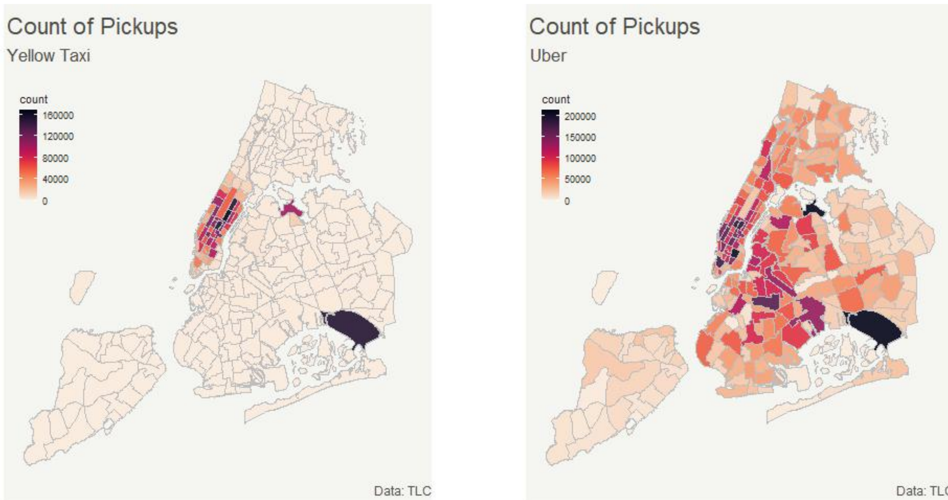

(a) Yellow taxi (b) Uber

Figure A1. : Pickups by Zone

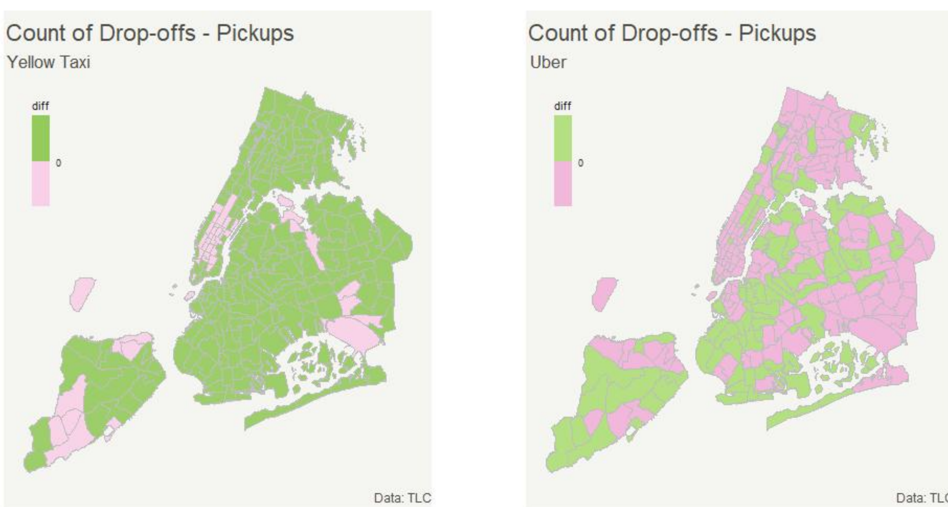

(a) Yellow Taxi (b) Uber

Figure A2. : Indicator of Dropoffs vs Pickups by Zone

Note: Since this is an indicator of drop-offs - pickups, the green zones indicate more drop-offs (than pickups) and the red ones show more pickups (than drop-offs) in a zone.

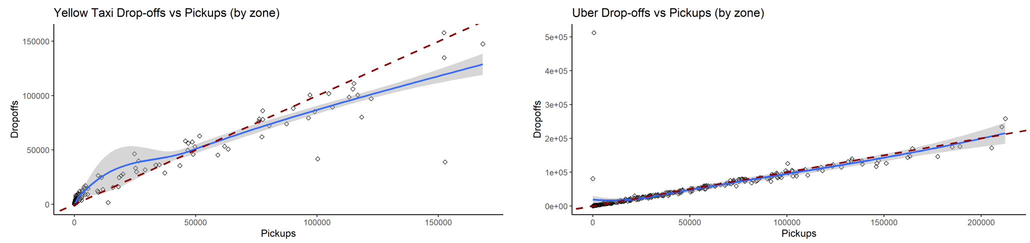

(a) Yellow Taxi (b) Uber

Figure A3. : Dropoffs vs Pickups Scatter Plots

Note: The number of pickups in a zone are on the x-axis and the number of drop-offs on the y-axis. The dashed red line is at 45 degrees.

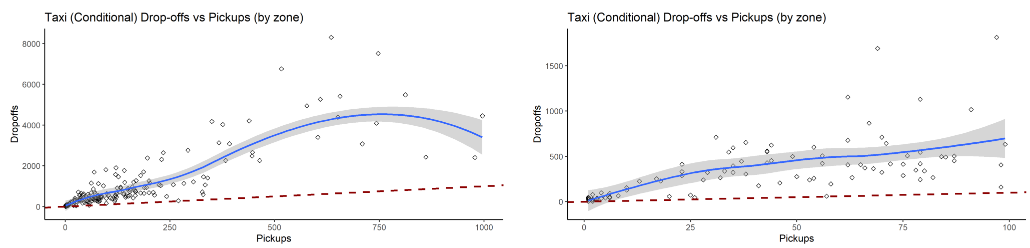

(a) Zones with < 1,000 pickups (b) Zones with < 100 pickups

Figure A4. : Dropoffs vs Pickups Scatter Plots (Subset by pickups)

Note: This is zooming into the bottom left part (low number of pickups and drop-offs) of the yellow taxi graph from figure A3.

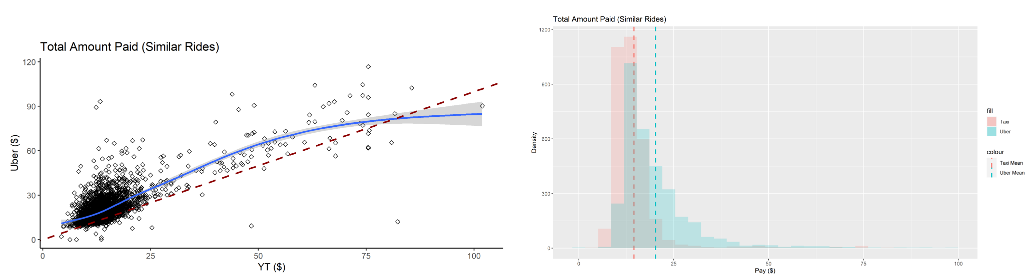

(a) Scatter plot (b) Distribution

Figure A5. : Matched trips: Uber price vs yellow taxi price

Note: I use data on matched trips to compare the total amount paid by passengers for the two cab types. The x-axis is ride fare for yellow taxi, and y-axis is for Uber.

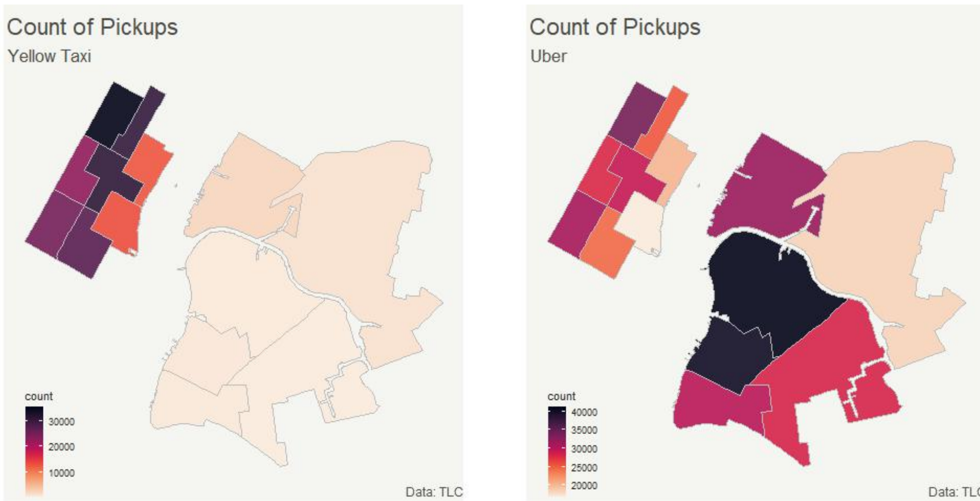

(a) Yellow Taxi (b) Uber

Figure A6. : Pickups by Zone (Subset)

Note: The top left are zones from lower Manhattan, while those on the bottom right are in the boroughs (at the edge of Brooklyn and Queens). They're connected by the Queens Midtown Tunnel.

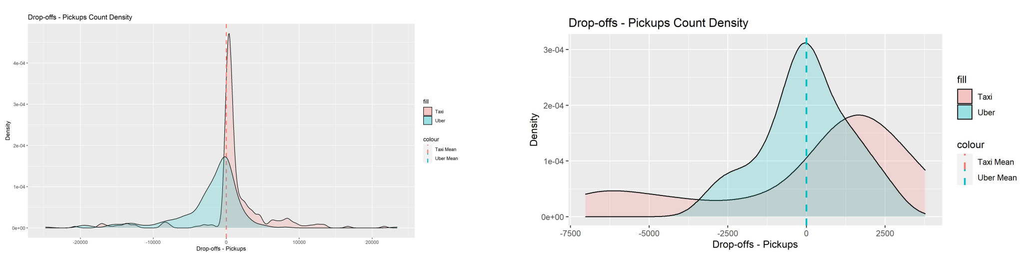

(a) All Zones (b) Subset zones

Figure A7. : Drop-offs - Pickups Density

Note: The left panel plots a distribution of drop-offs - pickups for all the zones in the data. The right panel plots the same for only the subset of zones considered in this paper. Here, red is 'yellow taxi' and blue is 'Uber'. A distribution skewed to the right implies that there are more zones where drop-offs > pickups.

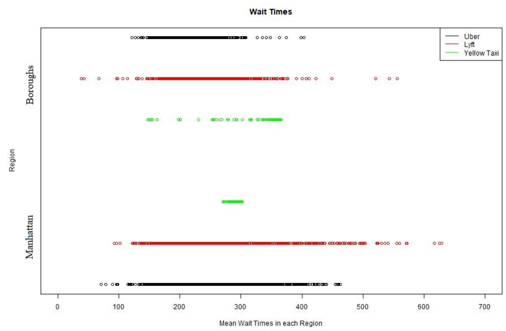

Figure A8. : Wait Times by Region and Cab type

Note: This figures plots the observed (for Uber and Lyft) and simulated wait times (for yellow taxi) in the regions considered for demand estimation.

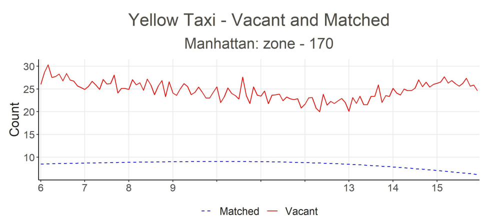

(a) Yellow Taxi

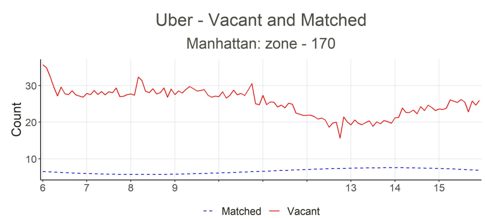

(b) Uber

Figure A9. : Vacant Cab Distribution - Manhattan

Note: These show the vacant cab distribution compared to observed matches in the same Manhattan zone.

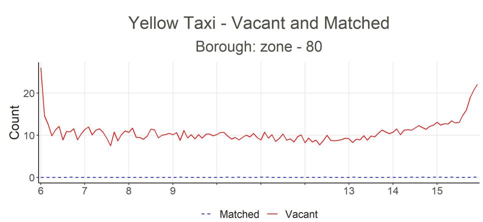

(a) Yellow Taxi

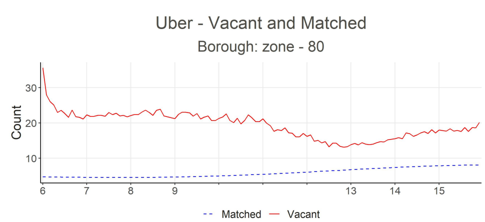

(b) Uber

Figure A10. : Vacant Cab Distribution - Boroughs

Note: These show the vacant cab distribution compared to observed matches in the same borough zone.

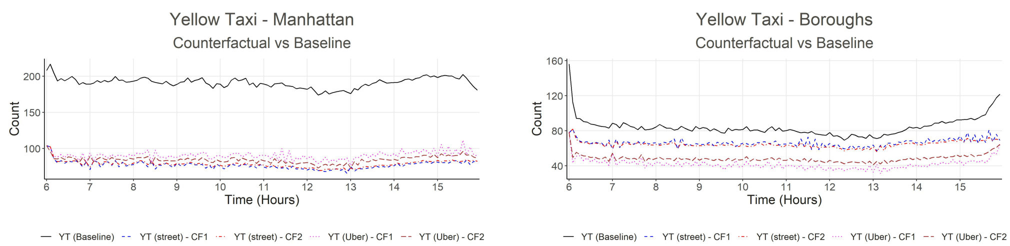

(a) Manhattan (b) Boroughs

Figure A11. : Aggregate Counterfactual Vacant Cab Distribution for yellow taxi

Note: These show the total vacant YT distribution in Manhattan and Boroughs over time for the two counterfactuals (CF 1: 'Uber react'; CF 2: 'Demand change'). 'Street' is when a cab searches for street-hailing, and 'Uber' is when it switches to the app.

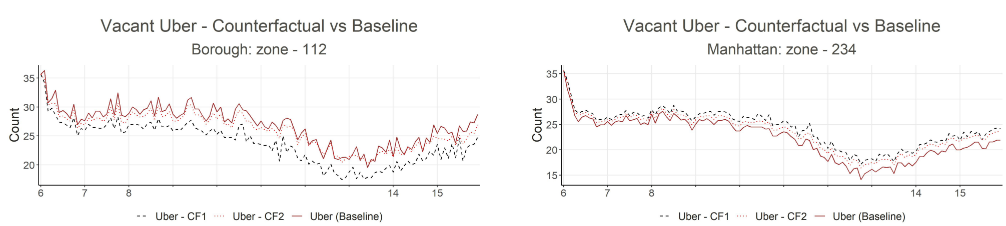

(a) Manhattan (b) Boroughs

Figure A12. : Counterfactual Vacant Cab Distribution for Ubers

Note: These show the vacant Uber distribution for a Manhattan and a Borough zone.