

Ride-share matching algorithms generate income inequality

Eszter Bokányi^{a,b} and Anikó Hannák^{* c,d,e}

^a*Eötvös Loránd University, Budapest, Hungary*

^b*Agglomeration and Social Networks Lendület Research Group, Centre for Economic and Regional Studies of the Hungarian Academy of Sciences*

^c*Complexity Science Hub, Vienna, Austria*

^d*Vienna University of Economics and Business, Austria*

^e*Centre for Economic and Regional Studies of the Hungarian Academy of Sciences*

Abstract

Despite the potential of online sharing economy platforms such as Uber, Lyft, or Foodora to democratize the labor market, these services are often accused of fostering unfair working conditions and low wages. These problems have been recognized by researchers and regulators but the size and complexity of these socio-technical systems, combined with the lack of transparency about algorithmic practices, makes it difficult to understand system dynamics and large-scale behavior. This paper combines approaches from complex systems and algorithmic fairness to investigate the effect of algorithm design decisions on wage inequality in ride-hailing markets. We first present a computational model that includes conditions about locations of drivers and passengers, traffic, the layout of the city, and the algorithm that matches requests with drivers. We calibrate the model with parameters derived from empirical data. Our simulations show that small changes in the system parameters can cause large deviations in the income distributions of drivers, leading to a highly unpredictable system which often distributes vastly different incomes to identically performing drivers. As suggested by recent studies about feedback loops in algorithmic systems, these initial income differences can result in enforced and long-term wage gaps.

As they grow in popularity, ride-sharing and food-delivery services such as Uber, Lyft, Ola or Foodora are quickly transforming urban transportation eco-systems [1, 2]. These services have revolutionized most aspects of the transportation market. By managing the rides through a mobile application, they lower the entry barriers to the service for both users or passengers and drivers. The rating system facilitates trust between drivers and users, and the flexible working hours make ride-sharing services a popular choice for people starting a new career or a side-job.

A key feature of these services is that an algorithm replaces human dispatchers in the task of matching available drivers to the incoming requests. Companies are now able to optimize the matching

^{*}To whom correspondence should be addressed. E-mail: aniko.hannak@wu.ac.at

with unprecedented precision using data they possess on cars, drivers, and traffic conditions [3], resulting in better service availability, shorter waiting times, and ultimately a boost in company profits [4]. On the other hand, in the process of maximizing profit, drivers’ interests get sidelined, and undesirable social outcomes can emerge [5–8].

Recent studies raise concerns about the risks threatening workers well-being, including racial bias, worker safety, fairness to workers, and asymmetries of information and power [9–13]. Since workers are not able to obtain remedies through official channels [14, 15], strikes have gotten increasingly common in the past years. Drivers of Uber, Lyft, Ola, Foodora demand higher fares, job security, and livable incomes all over the world [16].

While traditional taxi services allowed all workers to hear the same information and talk to the dispatcher, in these modern systems drivers see nothing besides their next target [15]. The automated algorithmic features act as a barrier between employees and management, annihilating this relationship and ultimately preventing worker feedback to be incorporated into the design. Moreover, the proprietary nature of these systems, there is limited access to data or the rules of the system [5, 17, 18]. Both regulators and researchers struggle to gain a deeper understanding of the ongoing problems and the underlying processes, which ultimately hinders informed intervention or monitoring by regulatory agencies [19, 20].

Most existing literature in the area of taxi matching algorithms is concerned with optimizing aggregate outcomes for the whole system [21–27]. This approach aims to maximize the benefits for the company or to minimize the adverse effects such as CO₂ emissions, overall distances driven or the passenger waiting times. Following the line of fairness measurement literature [28–31], we instead focus on the fair distribution of income from the drivers perspective, because current systems do not guarantee the same income for the same amount of work, neither across workers nor over time [10–12, 32].

We use an agent-based simulation to systematically study the mechanisms in ride-hailing and delivery systems from the perspective of the drivers. We first quantify the income inequality level of systems that use a company-level profit maximizing approach, and explore how the system-level behavior changes as a function of input parameters like the number and distribution of taxis and passengers, city layout, and driving strategies. Next, we investigate the trade-off between fairness of driver incomes and loss in overall revenue through an algorithm, designed to integrate the fairness perspective into the matching of drivers with requests.

1 Income Inequality in Ride Hailing

We create a simulated city environment with drivers, passengers and various parameters which we will introduce throughout the section. Our goal is to simulate a diverse range of cities and real-world traffic conditions, thus we derive our model parameters from real-world data sets.

Taxis in our city drive along a grid, moving one block at each time step. In our basic setting both request and drop-off locations are most likely to be in the city center, in line with previous studies on real-world data [33–35]. The matching algorithm between drivers and requests is similar to the algorithm that Uber and most taxi companies use: passengers are matched with the closest available car. Additional parameters capture various real-world scenarios such as city layout, changes in supply and demand, driver strategies, and different settings of the matching algorithm. The pricing scheme is similar to that of UberX in Boston [36], and fuel costs are accounted for [37]. We run all simulations for what equals a 40-hour work week (see 3 for details).

As described before, major complaints of drivers are low wages, and the unpredictability and unfair distribution of incomes. Low wages could only be addressed by the company itself, because they are in control of the payment system. Thus, we focus on understanding the two other complaints through investigating the distribution of incomes across workers for various settings of our system.

1.1 Changes in Supply and Demand

We use two variables to calibrate the supply and the demand in our system: the supply d (density) is defined as the number of taxis per square kilometer, while R captures demand-to-supply ratio, which is the fraction of the demanded travel distance over the supplied travel distance (see 3 for details). To pick parameter spaces for d and R that cover actual real-world scenarios, we calculate the number of cars and passengers using empirical data from the NYC Taxi and Limousine Commission [38].

Figure 1A depicts each hour of year 2013 in NYC in the $d - R$ parameter space. Plotting an example day, 15 January 2013, from the same dataset with its hourly parameter values shows, that even within one day and one city, both parameters strongly vary. While such detailed data is not available from other cities, we analyzed aggregate statistics as well as extreme events such as public transportation failures, strikes or bad weather conditions (see Section 3). We conclude that the taxi density range $d \in [3, 30]$ [1/km²] and the demand-to-supply ratio range $R \in [0.1, 1]$ will capture a wide range of city sizes, seasonal changes, varying traffic at different times of the day, and even extraordinary events that cause sudden changes in the demand or supply.

Figure 1B shows how the average income changes with growing number of requests, but constant taxi numbers, that is, a constant density d . Each marker denotes the averaged result of 10 simulations, ran with the associated parameters. As expected, the income is directly proportional to the demand,

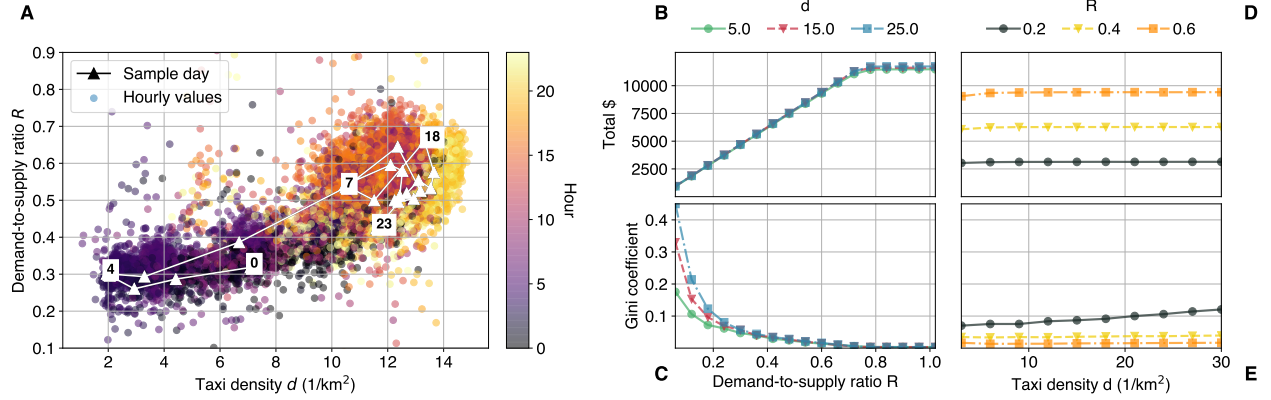


Figure 1: (A) Hourly values of the taxi density d and the demand-to-supply ratio R throughout year 2013 in NYC. Parameter changes within one selected day, 15 January are plotted as white connected triangles. The total income (B) and Gini coefficient (C) for varying demand to supply ratio, at three fixed density values. We see that income grows linearly as the demand to supply ratio grows. Higher demand leads to decrease in inequality, for all density levels. The total income (D) and Gini coefficient (E) for varying density values, at three fixed demand to supply ratio values. Higher demand generates higher incomes. Inequality grows with more cars in a city.

regardless of d . For example by doubling the demand-to-supply ratio R from 0.3 to 0.6, the income roughly increases from \$4700 to \$9400 for all three d values. Thus, as long as the system is able to serve all the passengers, demand determines the total income. After a certain point, incomes saturate as taxis are not able to serve all requests, and the system reaches its maximum capacity around $R = 0.78$.¹

Figure 1C measures the inequality for the same parameters using the Gini coefficient of the incomes at the end of the simulation. Gini coefficient is an inequality measure that captures the deviation of the Lorenz curve of the income distribution from that of an ideal one, where a given cumulative percentile of the population holds the same percentile amount from the incomes (see Section 3). For low demand, that is, low R , the Gini coefficient starts at high values, with 0.18 for $d = 5$ [1/km²], 0.32 for $d = 15$ [1/km²] and 0.45 for $d = 25$ [1/km²]. As the demand increases, the Gini coefficient decreases and converges for different taxi densities, with the Gini at $R = 0.6$ equalling to as little as 0.01 for all three measured d values.

Figures 1D-E depicts the effect of increasing traffic (more demand and more supply), that is, constant R , but increasing taxi density d . Again, higher demand generates higher average incomes, that is roughly \$3100 for $R = 0.2$, \$6200 for $R = 0.4$, and \$9400 for $R = 0.6$. The income is only very slightly affected by the taxi density in a city, as it increases by only 3% for a tenfold increase of $d = 3$ [1/km²] to $d = 30$ [1/km²]. However, higher density leads to a more unequal distribution

¹ We caution the reader that fares highly vary across cities, service providers, and over time even within the same company. While we present the incomes in \$, the values are approximations of the fares and the emphasis is on the relative value and the shape of the distributions.

of those same average incomes, as seen from Figure 1E. For $R = 0.2$, the Gini coefficient almost doubles with the same a tenfold increase of taxi density $d = 3$ [1/km²] to $d = 30$ [1/km²] that caused the 3% increase in the incomes. For higher demand, this inequality shift is smaller: while there is a 14% increase with changing $d = 3$ [1/km²] to $d = 30$ [1/km²] for $R = 0.4$, the Gini is almost constant for $R = 0.6$.

We note, that the Gini coefficients for the high supply (high d) and low demand-to-supply (low R) regions are comparable to the Gini coefficients observed for country-wise incomes [39]. Given that national statistics encompass inequalities across different professions and sectors, and we only compare workers of the same sector, with the same working hours and even with the same fitness, these values translate to a surprising level of inequality.

1.2 Varying Spatial Activity

Next, we consider variations in the city layout and the traffic flow patterns. Apart from the simplest city-center scenario, we run simulations for i.) a city with a larger center, ii.) a city with 2 centers, and iii.) cases where the pickup and destination locations do not to overlap (see the detailed description of these layouts in Section 3). Our motivation is explained by Figure 2A which shows the temporal and spatial changes in the difference between trip origin and destination location distributions in NYC. We can clearly see the average flow of passengers in the morning towards the center for work (Figure 2B) and away from the center in the evening (Figure 2C).

Before we turn our attention to income inequality in the above layouts, it is worth investigating whether and how they affect the overall income, since they all share the same R given identical traffic, i.e., the same demand-to-supply ratio. Figure 2D shows that the average income strongly depends on the distribution of spatial activity, when examining a setting with fixed $d = 15$ [1/km²] and $R = 0.48$. If the distribution of origins and destinations overlap completely, more income is generated regardless of the details, e.g. the spreading of the city center or having multiple centers. On the other hand, in the cases where there is a dominating flow of vehicles due to asymmetries in the origin and destination demands, the income is lowered to roughly 50% (Outwards flow) and 20% (Inwards flow). Moreover, the inequality is significantly higher over the whole demand-to-supply R range (Figure 2E). In the Inwards flow case, the Gini coefficient starts from as high as 0.76, but decreases fast thereafter, and reaches 0.07 for $R = 0.6$. When the traffic flows mostly outwards from the city-center, the initial Gini of 0.45 decreases very slowly, though, and stays as high as 0.38 for $R = 0.6$. The above results indicate that the distribution of spatial activity strongly influence both the average income and the income inequality.

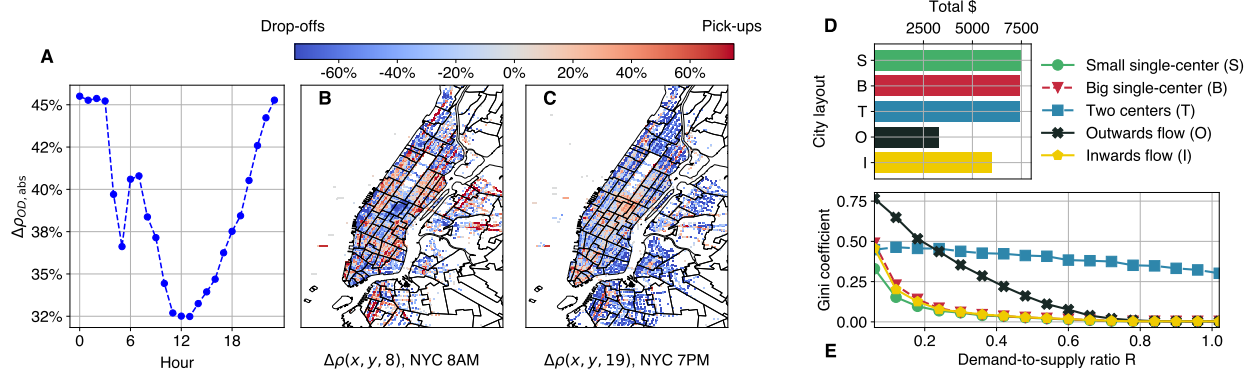


Figure 2: (A-C) Temporal and spatial changes in the difference between trip origin and destination locations in NYC. (A) Hourly average of the unsigned symmetric difference $\Delta\rho_{OD,abs}$ in NYC pixel areas. The figure indicates that origin and destination patterns vary throughout one day. (B)-(C) Spatial distribution of ρ_{OD} at 8AM and 7PM in NYC. Red values indicate more pickups, blue values more drop-offs in a pixel area. At 8AM, there is an influx into the center of Manhattan and other business and work areas. At 7PM, the direction of the taxi flow is reversed, as people travel towards the outskirts from the center. See Section 3 for details on $\Delta\rho_{OD,abs}$ and ρ_{OD} . (D) Average income of drivers depending on city layout. Generally, the overlapping cases generate more income, and movement towards the city center as well. (E) Gini coefficient of the seven layouts with varying demand to supply ratios. As demand grows, inequality decreases for every layout. Non-overlapping layouts (flows) in the city generate higher inequality.

1.3 Waiting vs Cruising

An important decision every driver faces many times throughout the day is what to do in the time period until the next request comes in, which might amount to a considerable share from the overall time spent online [3, 20]. So far in our simulations, the cars were waiting in place after dropping off their passengers. Another obvious idling strategy is to cruise back towards the center to meet more demand. While in a real system, drivers presumably use a mix of these two strategies, here, we investigate the outcome of the two extreme cases, namely when all taxi drivers either wait or head towards the city center.

Intuitively, we would expect that the cruising strategy leads to similar throughput in the cases of overlapping pick-up and drop-off distributions, smaller overall throughput in the cases where traffic flows towards the center, and higher overall throughput when population flows towards the outskirts. But does higher throughput also lead to lower Gini coefficients?

Figure 3A-E shows the waiting/cruising scenarios side by side for $d = 15$ [$1/\text{km}^2$] and $R = 0.4$. We observe significant differences in fairness between the waiting and cruising strategies in all city layouts. In the case of overlapping pick-up/drop-off location distributions the waiting strategy is fairer, as illustrated by the narrower distributions on Figure 3A-C corresponding to lower Gini coefficients while the average income is untouched. In case of the asymmetrical layouts, the strategy

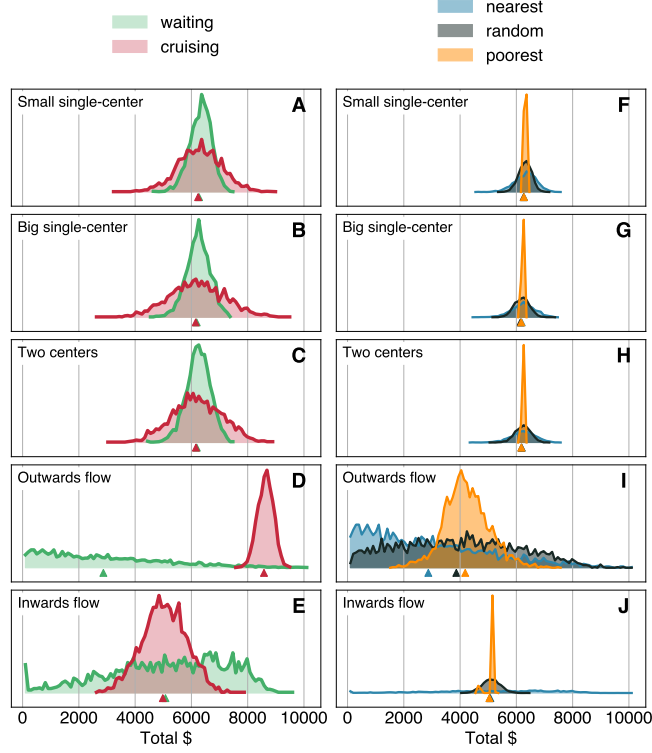


Figure 3: (A-E) Distribution of incomes of different cruising strategies for different city layouts at $R = 0.4$ and $d = 15[1/\text{km}^2]$. In symmetrical layouts the cruising strategy leads to more unequal distribution of incomes with similar means. Among the asymmetrical layouts, in the case of outwards flows, cruising towards the center leads to higher incomes. (F-J) Distribution of incomes of different matching algorithms in different city layouts and with the waiting strategy, at fixed $d = 15 [1/\text{km}^2]$ and $R = 0.4$. The poorest matching strategy creates a more equal distribution for all of the investigated geometries. (A-J) The vertical scale of the distributions is omitted for better readability, but all distributions are normalized such that the areas under the curve are equal to 1. Triangles mark the means below the distribution curves.

of cruising back to the center increases fairness, and in the Outwards flow layout it even raises the average income (decrease of the Gini from 0.22 to 0.07 and average income increases by almost 200%).

These results underline the importance of transparency and the direct effect of information asymmetry on drivers, who in the current setup of ride-hailing systems can not make informed decisions about their strategies. Moreover, it shows that a seemingly small changes in the system settings can lead to large differences in the fairness guarantees of the overall system.

1.4 Fixing the algorithm

Lastly, we examine whether we can incorporate the fairness perspective into our system and achieve more equal incomes while conserving the overall revenue, similarly to the perspective of [29]. Our goal is to keep track of drivers’ income throughout the day and take it into account when assigning rides.

With this idea in mind, we create the *poorest* algorithm: a modification of the current matching algorithm which keeps track of drivers’ income at each point in time and assigns taxis not only based on distance from the potential passenger but also the money they made so far (see details in Section 3). To meaningfully compare the algorithms, we include a baseline algorithm which assigns drivers to passengers randomly within a given radius. This *random* setup should create higher fairness but lower total income than picking the nearest available vehicle.

Figure 3F-J shows the income distribution for the three algorithms for all city layouts. The narrow distributions of the *poorest algorithm* in the symmetric cases show that this poorest correction effectively mitigates the adverse effects of the *nearest algorithm*. In these cases, *poorest* performs even better than the *random assignment* that we consider as a baseline for fair conditions. Moreover, the strategy also helps mitigating inequalities on the Inwards flow layout. While not as strongly as with the other layouts, poorest significantly increases fairness with the Outwards flow layout, and even increases the mean income. Since the nearest algorithm mostly assigns drivers from the center, more and more drivers end up in the outskirts without close-by rides. As the poorest algorithm is more likely to pick a driver stranded on the outskirts, it compensates this undesired process and ultimately leads to higher mean income.

2 Concluding Discussion

We presented results of an agent-based automated taxi system simulation from a so-far unexplored angle, namely, the fairness of incomes among drivers. Our simulation environment allows us to cover a wide spectrum of city layouts and traffic conditions, and to test arbitrary matching algorithms. We find that a low demand-to-supply ratio (more taxis/less passenger requests) leads to greater inequality, and that the inequality is largely dependent on the spatial distribution of request origins and destinations, on the chosen idling strategy of the drivers, and the matching algorithm.

We also proposed a new matching algorithm that attempts to equalize incomes by promoting drivers on the lower end of the current income distribution. This method significantly boosts fairness in driver incomes in all investigated scenarios, while not impairing, and sometimes even increasing average incomes through stimulating optimal spatial redistribution of the taxis.

Notably, our most surprising result is that drivers with the same qualifications and working hours can end up with vastly different incomes by chance. One might argue that these income differences, observed over a short time period, disappear on the long run. This would indeed be the case under our model if used for long-term predictions (even if the regression to the mean could take a considerable amount of time). However, a large body of sociological and economic literature has shown that emerging inequalities are amplified through feedback loops and other processes that are not directly captured by the simulation in this paper [13, 40]. The initial differences combined with emerging feedback loops and varying skills, working hours, and ratings can result in increased and continuing wage gaps. Moreover, literature shows that workers have to focus on daily and weekly targets because of regular fees and payments [41]. Thus, short-term income variance creates unpredictable circumstances and may cause drivers to take risks such as unsafe amounts of overtime [15].

Due to the sheer size and complexity of the problem, our study is limited in different ways. First, the city model generalizes to a square grid street network, and the traffic model averages over highly dynamic conditions such as the true taxi velocity, which might depend on the time of day, road type, and physical obstacles. For example, Uber’s surge pricing, which aims to steer supply and demand, is an important factor in driver earnings and another source of unpredictability. In contrast, our algorithms are only elementary dispatching systems that are likely not identical to the proprietary solutions.

Impact Our results underline the urgency to establish the necessary means for drivers to make informed decisions while working. In our view, giving more control to workers is only helpful when coupled with insights that help them make informed decisions. Currently, drivers’ basic decisions, such as when and where to work, are based on limited and local information. Our systematic analysis aims to fill this gap by connecting local knowledge with global outcomes while focusing on drivers’ interests. We derive results that directly help drivers align short-term strategies with their long-term goals based on a few key system indicators such as demand, supply or the spatial distribution of requests.

Our work will be the first to take fairness metrics into account in the evaluation of system performance. Currently, companies focus on maximizing their own profit, which does not align with worker welfare and fairness. Technical research has followed the former demand, considering only optimization problems of system level efficiency. By sharing our simulation system on GitHub², we hope to encourage the research community to work towards a thorough understanding of the effects of algorithm choice on social outcomes, and that eventually service designers will be provided with specific technical recommendations on how to implement a possible fairness perspective. Ultimately, fairness guarantees, either provided by the platforms themselves or via external monitoring of policy

²<https://github.com/bokae/taxi>

agencies, are necessary to increase the trust between ride-hailing companies and their workers.

Lastly, we want to emphasize the advantages of complex systems methods in algorithmic fairness research. Currently, most algorithmic auditing research projects focus on identifying issues with individual platforms at particular points in time [42, 43]. We presented an enhanced approach to algorithmic auditing that takes into account system dynamics, complex phenomena, and downstream consequences of algorithms. This allows the generalization of insights across platforms, pointing towards fundamental properties of socio-technical systems. Complex algorithmic auditing has the potential to be adapted to research harmful phenomena beyond this project, such as the filter bubble effect, discrimination, or polarization in algorithmically aided systems [44–46].

3 Materials and Methods

Simulated city We investigate generalized and simplified ride-sharing systems with an agent-based simulation. In this simulation, the city is represented as a square grid similarly to [47], consisting of n times ($n = 20$) m pixels ($m = 20$). To associate simulation units to real-world dimensions, the side length of the pixels is set to $du = 100$ m in SI distance units.

Time, movement and distances In each simulation time step, taxis move exactly one distance unit on the grid, and they keep this constant average velocity v throughout the whole simulation. We set v to be 10 m/s, a hypothetical average daytime speed as in [21], which implicitly defines simulation time unit: $tu = a/v = 10$ s. All presented results throughout the paper correspond to simulations ran for what equals a 40-hour work week, that is, until $5 \cdot 8 \cdot 3600/tu = 14400$ simulation time units. Literature suggests one work week to be a reasonable time frame for which workers might expect predictable and stable incomes because of weekly fixed costs [41].

Supply and demand There are T taxis in the system. We control the number of taxis through a density parameter d , that is the number of taxis per total area of the system. The dimension of d is $[1/\text{km}^2]$.

$$d = \frac{T}{A} = \frac{T}{n \cdot m \cdot du^2}. \quad (1)$$

The demand-to-supply ratio is characterized by a dimensionless measure R , which is the fraction of the total demanded travel distance over the distance provided by the taxis during the same time window:

$$R = \frac{\lambda \cdot l \cdot t}{T \cdot v \cdot t} = \frac{\lambda \cdot l}{T \cdot v}, \quad (2)$$

where l is the average trip length requested by passengers in the system, t is the elapsed time, and λ is the average number of requests per time unit.

Traffic flow and city layout Since the spatial distribution of the demand is typically not uniform in a city [33–35], we introduce different probability distributions that characterize the request origins and destinations. In the simplest case, the origin and destination distributions are identical 2D Gaussians, mimicking scenarios where demand is higher in the city center, and gradually decreases further from the core [48, 49]. The scenarios where origin and destination distributions are not overlapping, such as in the case of morning rush hours or late night rides to residential areas, are simulated by defining separate Gaussian distributions for request origin and destination, respectively, with different standard deviations or centers. Table 1 provides the different distributions in our simulations, with f being the weight of Gaussian mixtures, μ their centers, and σ their standard deviations. We compared the distribution of trip lengths generated from these Gaussian distributions to that of a real-world dataset [50]. We concluded that the distribution from a real day is the composition of many different system states and spatial activities, therefore, we do not require a strict match between the statistics of the two. We also generated layouts with different cutoffs (exponential, sigmoid) in the decrease of the density of the pick-ups and drop-offs towards the outskirts, but the results did not differ qualitatively from those obtained with Gaussian distributions.

	Origin			Destination		
	μ	σ	f	μ	σ	f
Small city-center	[20,20]	10	1.0	—		
Large city-center	[20,20]	20	1.0	—		
Multi-centered city	[12,12]	8	0.5	—		
	[28,28]	8	0.5			
Outwards flow	[20,20]	6	1.0	[20,20]	12	1.0
Inwards flow	[20,20]	12	1.0	[20,20]	6	1.0

Table 1: List of different origin and destination distributions in the simulations. Each three-column row represents a 2D Gaussian distribution with a mean at μ in simulation coordinates, a standard deviation σ , and a weight f . Distributions having more than one row are a mixture of the listed Gaussians with a weight f . The — symbol denotes that the destination distribution is the identical to the origin distribution.

Implementation Once we have the taxis and incoming requests in the system, we have to define how these elements behave while running the simulation. If a passenger makes a request, the origin and destination location is submitted to the waiting queue, and the request is first marked as pending. At each time step, we try to assign an available taxi to all requests in the queue, from oldest to newest. If the assignment is successful, the request is marked as confirmed, and the assigned taxi is marked as booked, and it ceases to be available. When the taxi reaches the request origin, it picks up the passenger, and the taxi and the request both enter into serving state. A

path is created for the taxi, along which it moves towards the request destination as the time is running. Once the destination is reached, the taxi drops off the passenger, the request is marked as completed, and the taxi becomes available again. Requests that pend too long, that is, for more than 30 simulation units (corresponding to 5 min in real time units) in the pending queue are canceled, as passengers tend to not waiting very long.

Income Because the pricing scheme might differ between different service providers, cities, or it might change rapidly even within the same system (similar to surge pricing in Uber [51]), we only introduce illustrative USD values based on real UberX fares from an example US city, Boston [36]. Thus, the income has a fixed \$2/trip and a distance-based \$1/km part, that already includes the mile-based and minute-based fees calculated using an average velocity of 10 m/s. Note that the final income values only serve as indicators of general income trends in mean incomes and inequalities, and they shall not be interpreted as realistic incomes. Fixed costs are assumed to be identical for each driver, thus they are omitted from the calculations. Fuel consumption is proportional to the total distance covered (going for and serving a request, cruising while empty). With an average price of fuel of \$3 per gallon [37], an estimated 10 l per 100 km [31] consumption leads to $\approx \$0.008/du$ costs.

Matching algorithm The essence of the assignments between taxis and requests is the matching algorithm. All presented algorithms first consider a pool of available drivers in a 1 km range around the request origin. First, we implement the *nearest* algorithm, that looks for the nearest available taxi from the pool. This algorithm is the basis for many existing solutions [25, 27, 52, 53]. We also propose a mechanism that takes the drivers’ previous incomes into account and the driver with the least income so far (*poorest*). As a baseline, we also include a method with uniform random assignment from the pool (*random*).

We can also set different initial conditions and different idle behaviour for the taxi drivers. They either start the day from a "taxi base" in the center location, or from a random location on the grid that mimics their "home" location. After a completed ride, they can either be *waiting* for their next ride in place, or they can start *cruising* back towards the city center. There is also a possibility for simulating a realistic 8-hour shift, where after the end of the shift, every taxi completes its last assignment (if it has one), and then starts the new shift again from their "home/taxi base".

The code for the simulation is available at <https://github.com/bokae/taxi>.

Supply and demand parameter ranges To match our simulation to real-world parameter values, we obtained the only comprehensive ride-sharing vehicle dataset that is available for the vehicles operating in New York City from year 2013 [50]. We calculate average trip length l ,

average velocity v , number of distinct taxis operating T , and number of served requests per second λ for each hour of the year from the dataset. Using equations (1) and (2), we calculate the two main system parameters d and R for these hours using $A = 783 \text{ km}^2$ as the area for NYC. Figure 1A shows the calculated values, and the changes in d and R within one sample day on 15 January 2013. According to the figure, d varies between 2 and 15 $[1/\text{km}^2]$, and R between 0.2 and 0.75. However, we do not have data on the number of cancelled requests. Therefore, it might happen even within this dataset, that the real value of R is higher than 0.75 during peak hours. We also calculate approximations of d and R based on aggregated metrics for other towns. For example, we take a typical example for an oversupplied taxi system, the traditional taxi service in Berlin, with $d \approx 1 - 5 [1/\text{km}^2]$ throughout one day, and a typical $R = 0.5$ in peak hours [54]. An opposite scenario of undersupply is present in the city of Barcelona, where d varies between 15 $[1/\text{km}^2]$ and 35 $[1/\text{km}^2]$, and R between 0.15 and 0.5 [55]. Though representing very different real-life situations, in terms of overall numbers, both cities fall into our simulated parameter range. Extreme events such as public transportation failures or bad weather conditions might cause a sudden increase in demand, sometimes as high as doubling R [55]. Therefore, we conduct our simulations for up to $R = 1$.

City layouts We also investigate whether an expected shift in the direction of the trips happens during a day. First, we divide NYC into rectangular pixels of 0.01° longitude and 0.01° latitude, and we calculate the normalized distribution of pick-ups $\rho_O(x, y, h)$ and drop-offs $\rho_D(x, y, h)$ for each hour $h = 0, \dots, 23$ averaged over one year. We filter those x, y longitude-latitude pixels in the dataset, where the overall number of both pick-ups and drop-offs is below 5, and where $x \notin [-74.05, -73.9]$, and $y \notin [40.67, 40.85]$.

Normalization means that for each h ,

$$\sum_{x,y} \rho_O(x, y, h) = \sum_{x,y} \rho_D(x, y, h) = 1. \quad (3)$$

We then define the symmetric difference between these two distributions:

$$\Delta\rho_{OD}(x, y, h) = \frac{\rho_O(x, y, h) - \rho_D(x, y, h)}{\rho_O(x, y, h) + \rho_D(x, y, h)} \quad (4)$$

Figure 2A shows the daily mean of the absolute values of the symmetric differences in the pixels

$$\Delta\rho_{OD,\text{abs}}(h) = \frac{1}{N} \sum_{x,y} |\Delta\rho_{OD}(x, y, h)|, \quad (5)$$

where N is the number of x, y pairs in the analysis. Figure 2B and C show $\Delta\rho_{OD}(x, y, 8)$ and $\Delta\rho_{OD}(x, y, 19)$, respectively, on the taxi zone map of NYC. We can observe the inflow of people

into the city center in the morning, and the outflow of the traffic around 7PM.

Fairness metric and its interpretation To assess the inequality in the income distributions of the drivers after 1 week’s worth of work, we calculate the Gini coefficient [56]. The Gini coefficient is an important tool in economics for analyzing income distributions over countries or regions. Its typical value for countries (being between 0.2-0.4) is summarized e.g. in OECD reports [39]. The value of the Gini coefficient is independent of the mean of the distribution, therefore, differences in the inequalities can be compared for scenarios generating different mean incomes in the simulation.

Acknowledgements

We would like to thank the helpful comments and insights of Johannes Wachs, Christoph Stadtfeld, Sándor Juhász, Gábor Hannák, Kenneth Joseph, Piotr Sapiezynski and David Garcia. During this research, Eszter Bokányi was supported by the ÚNKP-18-3 New National Excellence Program of the Hungarian Ministry of Human Capacities.

References

- [1] Brad Stone. Uber Is Winning Over Americans’ Expense Accounts, 2015. URL <https://www.bloomberg.com/news/articles/2015-04-07/uber-is-winning-over-americans-expense-accounts>. [Accessed May 28, 2019].
- [2] Jonathan V. Hall and Alan B. Krueger. An Analysis of the Labor Market for Uber’s Driver-Partners in the United States. *ILR Review*, 71(3):705–732, may 2018. ISSN 0019-7939. doi: 10.1177/0019793917717222. URL <http://journals.sagepub.com/doi/10.1177/0019793917717222>.
- [3] Shih-Fen Cheng and Thi Duong Nguyen. TaxiSim: A Multiagent Simulation Platform for Evaluating Taxi Fleet Operations. In *2011 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology*, pages 14–21. IEEE, aug 2011. ISBN 978-1-4577-1373-6. doi: 10.1109/WI-IAT.2011.138. URL <http://ieeexplore.ieee.org/document/6040748/>.
- [4] Dan Kedmey. Uber’s Surge Pricing: A Quick and Dirty Guide, 2014. URL <http://time.com/3633469/uber-surge-pricing/>. [Accessed May 28, 2019].
- [5] Ben Popper. Uber kept new drivers off the road to encourage surge pricing and increase fares, 2014. URL <https://www.theverge.com/2014/2/26/5445210/in-san-diego-uber-kept-drivers-off-the-road-to-encourage-surge>. [Accessed May 28, 2019].
- [6] David Shepardson and Heather Somerville. Uber not criminally liable in fatal 2018 Arizona self-driving crash: prosecutors, 2019. URL <https://www.reuters.com/article/us-uber-crash-autonomous/uber-not-criminally-liable-in-fatal-2018-arizona-self-driving-crash-prosecutors-idUSKCN1QM208>. [Accessed May 28, 2019].

- [7] Inae Oh. Uber Will Stop Charging Ridiculous Prices During Emergencies, 2014. URL https://www.huffingtonpost.com/2014/07/08/uber-surge-pricing{_}n{_}5568087.html. [Accessed May 28, 2019].
- [8] Alison Griswold. Uber price surging on Halloween: Does it take unfair advantage of drunk people?, 2014. URL http://www.slate.com/blogs/moneybox/2014/11/04/uber{_}price{_}surging{_}on{_}halloween{_}does{_}it{_}take{_}unfair{_}advantage{_}of{_}drunk{_}. [Accessed May 28, 2019].
- [9] Principles Fairwork, 2018. URL <https://fair.work/principles/>. [Accessed May 28, 2019].
- [10] Mark Graham and Joe Shaw. *Towards a fairer gig economy*. Meatspace Press, 2017. ISBN 9780995577664. URL https://issuu.com/meatspacepress/docs/towards{_}a{_}fairer{_}gig{_}economy.
- [11] Mark Graham and Jamie Woodcock. Towards a Fairer Platform Economy: Introducing the Fairwork Foundation. *Alternate Routes*, 29(Social Inequality and the Spectre of Social Justice): 242–253, 2018.
- [12] Min Kyung Lee, Daniel Kusbit, Evan Metsky, and Laura Dabbish. Working with Machines: The Impact of Algorithmic and Data-Driven Management on Human Workers. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15*, pages 1603–1612, New York, New York, USA, 2015. ACM Press. ISBN 9781450331456. doi: 10.1145/2702123.2702548. URL <http://dl.acm.org/citation.cfm?doid=2702123.2702548>.
- [13] Cody Cook, Rebecca Diamond, Jonathan Hall, John List, and Paul Oyer. The Gender Earnings Gap in the Gig Economy: Evidence from over a Million Rideshare Drivers. Technical report, National Bureau of Economic Research, Cambridge, MA, jun 2018. URL <http://www.nber.org/papers/w24732.pdf>.
- [14] Confessions of an Uber Driver Empty Promises, 2015. URL <http://uberconfession.tumblr.com/post/108296995865/empty-promises>. [Accessed May 28, 2019].
- [15] Alex Rosenblat and Luke Stark. Algorithmic Labor and Information Asymmetries: A Case Study of Uber’s Drivers. *International Journal of Communication*, 10, 2016. ISSN 1932-8036. doi: 10.2139/ssrn.2686227.
- [16] Megan Rose Dickey. Uber and Lyft drivers are striking ahead of Uber’s IPO — TechCrunch, 2019. URL <https://techcrunch.com/2019/05/06/uber-and-lyft-drivers-are-striking-ahead-of-ubers-ipo/?guccounter=1>. [Accessed May 28, 2019].
- [17] Liat Clark. Uber denies researchers’ ‘phantom cars’ map claim, 2015. URL <http://www.wired.co.uk/article/uber-cars-always-in-real-time>. [Accessed May 28, 2019].
- [18] Mark Wilson. Uber Fudges The Position Of Local Drivers, But They’ve Got A Pretty Good Reason Why. *FastCompany*. URL <https://www.fastcodesign.com/3049169/uber-fudges-the-position-of-local-drivers-but-theyve-got-a-pretty-good-reason-why>.
- [19] Le Chen, Alan Mislove, and Christo Wilson. Peeking Beneath the Hood of Uber. In *Proceedings of the 2015 ACM Conference on Internet Measurement Conference - IMC '15*, pages 495–508, New York, New York, USA, 2015. ACM Press. ISBN 9781450338486. doi: 10.1145/2815675.2815681. URL <http://dl.acm.org/citation.cfm?doid=2815675.2815681>.

- [20] Shan Jiang, Le Chen, Alan Mislove, and Christo Wilson. On ridesharing competition and accessibility: Evidence from uber, lyft, and taxi. In *Proceedings of the 2018 World Wide Web Conference, WWW '18*, pages 863–872, Republic and Canton of Geneva, Switzerland, 2018. International World Wide Web Conferences Steering Committee. ISBN 978-1-4503-5639-8. doi: 10.1145/3178876.3186134. URL <https://doi.org/10.1145/3178876.3186134>.
- [21] Michal Kümmel, Fritz Busch, and David Z.W. Wang. Taxi Dispatching and Stable Marriage. In *Procedia Computer Science*, volume 83, pages 163–170. Elsevier, jan 2016. doi: 10.1016/j.procs.2016.04.112. URL <https://www.sciencedirect.com/science/article/pii/S1877050916301351>.
- [22] R M Jorgensen, J Larsen, and K B Bergvinsdottir. Solving the dial-a-ride problem using genetic algorithms. *Journal of the Operational Research Society*, 58(10):1321–1331, oct 2007. ISSN 01605682. doi: 10.1057/palgrave.jors.2602287. URL <https://www.tandfonline.com/doi/full/10.1057/palgrave.jors.2602287>.
- [23] Aamena Alshamsi, Sherief Abdallah, and Iyad Rahwan. Multiagent self-organization for a taxi dispatch system. In *8th international conference on autonomous agents and multiagent systems*, pages 21–28, 2009.
- [24] William A. Bailey and Thomas D. Clark. A Simulation Analysis of Demand and Fleet Size effects on taxicab service rates. *Winter Simulation Conference*, pages 838–844, 1987. ISSN 02750708. URL http://delivery.acm.org/10.1145/320000/318705/p838-bailey.pdf?ip=157.181.168.12&id=318705&acc=ACTIVESERVICE&key=C8586E1EAA35B39B.A93B98384D428521.4D4702B0C3E38B35.4D4702B0C3E38B35&{}_{}_acm{}_{}_{}_{}_{}_6c9a834b1ccb20709a60f02c9462aeee.
- [25] Michał Maciejewski and Kai Nagel. Simulation and dynamic optimization of taxi services in MATSim. *Investment Management and Financial Innovations*, 5(4):7–24, 2014. ISSN 18129358. doi: 10.1287/. URL http://svn.vsp.tu-berlin.de/repos/public-svn/publications/vspwp/2013/13-05/2013-06-03{}_{}_Maciejewski{}_{}_Nagel.pdf.
- [26] Michał Maciejewski and Joschka Bischoff. Large-scale microscopic simulation of taxi services. In *Procedia Computer Science*, volume 52, pages 358–364, 2015. doi: 10.1016/j.procs.2015.05.107. URL https://ac.els-cdn.com/S1877050915009072/1-s2.0-S1877050915009072-main.pdf?{}_{}_tid=17be4ffe-1737-4569-91af-eb7b0bef05b2&{}_{}_acdnat=1525809470{}_{}_03d401fe6c7cef1e06c980ec8eb0c279.
- [27] Michał Maciejewski, Joschka Bischoff, and Kai Nagel. An Assignment-Based Approach to Efficient Real-Time City-Scale Taxi Dispatching. *IEEE Intelligent Systems*, 31:68–77, 2016. ISSN 15411672. doi: 10.1109/MIS.2016.2.
- [28] Ben Hutchinson and Margaret Mitchell. 50 Years of Test (Un)fairness. In *Proceedings of the Conference on Fairness, Accountability, and Transparency - FAT* '19*, pages 49–58, New York, New York, USA, 2019. ACM Press. ISBN 9781450361255. doi: 10.1145/3287560.3287600. URL http://arxiv.org/abs/1811.10104%{}_{}_0Ahttp://dx.doi.org/10.1145/3287560.3287600http://dl.acm.org/citation.cfm?doid=3287560.3287600.
- [29] Abhijnan Chakraborty, Asia J Biega, Aniko Hannak, and Krishna P Gummadi. Fair Sharing for Sharing Economy Platforms. *Fair Sharing for Sharing Economy Platforms Abhijnan*, (August):2–5, 2017. doi: 10.18122/B2BX2S. URL <https://doi.org/10.18122/B2BX2S>.

- [30] Meike Zehlike, Francesco Bonchi, Carlos Castillo, Sara Hajian, Mohamed Megahed, and Ricardo Baeza-Yates. FA*IR. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management - CIKM '17*, pages 1569–1578, New York, New York, USA, jun 2017. ACM Press. ISBN 9781450349185. doi: 10.1145/3132847.3132938. URL <http://arxiv.org/abs/1706.06368><http://dx.doi.org/10.1145/3132847.3132938><http://arxiv.org/abs/1706.06368><http://dx.doi.org/10.1145/3132847.3132938><http://dl.acm.org/citation.cfm?doid=3132847.3132938>.
- [31] Jiancheng Weng, Quan Liang, Guoliang Qiao, Zhihong Chen, and Jian Rong. Taxi fuel consumption and emissions estimation model based on the reconstruction of driving trajectory. *Advances in Mechanical Engineering*, 9(7):168781401770870, jul 2017. doi: 10.1177/1687814017708708. URL <http://journals.sagepub.com/doi/10.1177/1687814017708708>.
- [32] Christian Fieseler, Eliane Bucher, and Christian Pieter Hoffmann. *Unfairness by Design? The Perceived Fairness of Digital Labor on Crowdworking Platforms*, volume 156. Springer Netherlands, jun 2017. doi: 10.1007/s10551-017-3607-2. URL <http://link.springer.com/10.1007/s10551-017-3607-2>.
- [33] Xinwu Qian and Satish V. Ukkusuri. Spatial variation of the urban taxi ridership using GPS data. *Applied Geography*, 59:31–42, 2015. ISSN 01436228. doi: 10.1016/j.apgeog.2015.02.011. URL <http://dx.doi.org/10.1016/j.apgeog.2015.02.011>.
- [34] Chengbin Peng, Xiaogang Jin, Ka-Chun Wong, Meixia Shi, and Pietro Liò. Collective Human Mobility Pattern from Taxi Trips in Urban Area. *PLoS ONE*, 7(4):e34487, apr 2012. ISSN 1932-6203. doi: 10.1371/journal.pone.0034487. URL <https://dx.plos.org/10.1371/journal.pone.0034487>.
- [35] Yu Liu, Fahui Wang, Yu Xiao, and Song Gao. Urban land uses and traffic ‘source-sink areas’: Evidence from GPS-enabled taxi data in Shanghai. *Landscape and Urban Planning*, 106(1):73–87, 2012. ISSN 01692046. doi: 10.1016/j.landurbplan.2012.02.012. URL <http://dx.doi.org/10.1016/j.landurbplan.2012.02.012>.
- [36] Uber Boston — Uber Prices & Services — Drive In Boston, 2019. URL <http://www.alvia.com/uber-city/uber-boston-2/>. [Accessed May 28, 2019].
- [37] U.S. average gas prices by year 1990-2018 — Statista, 2019. URL <https://www.statista.com/statistics/204740/retail-price-of-gasoline-in-the-united-states-since-1990/>. [Accessed May 28, 2019].
- [38] NYC Taxi Trips by andresmh, 2017. URL <http://www.andresmh.com/nyctaxitrips/>. [Accessed May 28, 2019].
- [39] Leonardo Gasparini and Leopoldo Tornarolli. A review of the OECD Income Distribution Database. *The Journal of Economic Inequality*, 13(4):579–602, dec 2015. ISSN 1569-1721. doi: 10.1007/s10888-015-9299-x. URL <http://link.springer.com/10.1007/s10888-015-9299-x>.
- [40] Arnout Van de Rijt, Soong Moon Kang, Michael Restivo, and Akshay Patil. Field experiments of success-breeds-success dynamics. *Proceedings of the National Academy of Sciences*, 111(19):6934–6939, 2014.
- [41] Henry S. Farber. Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers. *Journal of Political Economy*, 113(1):46–82, feb 2005. ISSN 0022-3808. doi: 10.1086/426040. URL <https://www.journals.uchicago.edu/doi/10.1086/426040>.

- [42] Christian Sandvig, Kevin Hamilton, Karrie Karahalios, and Cedric Langbort. Auditing algorithms: Research methods for detecting discrimination on internet platforms. *Data and discrimination: converting critical concerns into productive inquiry*, 22, 2014.
- [43] Aniko Hannak, Piotr Sapiezynski, Arash Molavi Kakhki, Balachander Krishnamurthy, David Lazer, Alan Mislove, and Christo Wilson. Measuring personalization of web search. In *Proceedings of the 22nd international conference on World Wide Web*, pages 527–538. ACM, 2013.
- [44] Seth Flaxman, Sharad Goel, and Justin M Rao. Filter bubbles, echo chambers, and online news consumption. *Public opinion quarterly*, 80(S1):298–320, 2016.
- [45] Aniko Hannak, Claudia Wagner, David Garcia, Alan Mislove, Markus Strohmaier, and Christo Wilson. Bias in Online Freelance Marketplaces: Evidence from TaskRabbit and Fiverr. In *20th ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW 2017)*, Portland, OR, February 2017.
- [46] Michael D Conover, Jacob Ratkiewicz, Matthew Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. Political polarization on twitter. In *Fifth international AAAI conference on weblogs and social media*, 2011.
- [47] Saurav Ranjit, Apichon Witayangkurn, Masahiko Nagai, and Ryosuke Shibasaki. Agent-Based Modeling of Taxi Behavior Simulation with Probe Vehicle Data. *ISPRS International Journal of Geo-Information*, 7(5):177, 2018. ISSN 2220-9964. doi: 10.3390/ijgi7050177. URL <http://www.mdpi.com/2220-9964/7/5/177>.
- [48] Jing Yuan, Yu Zheng, Xing Xie, and Guangzhong Sun. Driving with knowledge from the physical world. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '11*, page 316, New York, New York, USA, 2011. ACM Press. ISBN 9781450308137. doi: 10.1145/2020408.2020462. URL <http://dl.acm.org/citation.cfm?doid=2020408.2020462>.
- [49] Yu Liu, Chaogui Kang, Song Gao, Yu Xiao, and Yuan Tian. Understanding intra-urban trip patterns from taxi trajectory data. *Journal of Geographical Systems*, 14(4):463–483, 2012. ISSN 14355930. doi: 10.1007/s10109-012-0166-z.
- [50] Chris Wong. FOILing NYC’s Taxi Trip Data Chris Whong, 2014. URL <https://chriswhong.com/open-data/foil{ }nyc{ }taxi/>. [Accessed May 28, 2019].
- [51] Nicholas Diakopoulos. How Uber surge pricing really works, 2015. URL <https://www.washingtonpost.com/news/wonk/wp/2015/04/17/how-uber-surge-pricing-really-works/?utm{ }term=.b16b8c1e885d>. [Accessed May 28, 2019].
- [52] Kevin Mark Novak and Travis Cordell Kalanick. System and method for dynamically adjusting prices for services, 2013. URL <https://patentimages.storage.googleapis.com/3a/a7/19/82903f8e9761cc/US20130246207A1.pdf>. [Accessed May 28, 2019].
- [53] Charly Robinson La Rocca and Jean-François Cordeau HEC Montréal. Heuristics for Electric Taxi Fleet Management at Teo Taxi. Available at: http://chairelogistique.hec.ca/wp-content/uploads/2017/11/infor_article.pdf (2017-11-02), 2017. URL <http://chairelogistique.hec.ca/wp-content/uploads/2017/11/infor{ }article.pdf>.

- [54] Joschka Bischoff, Michal Maciejewski, and Alexander Sohr. Analysis of Berlin’s taxi services by exploring GPS traces. In *2015 International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS)*, pages 209–215. IEEE, jun 2015. ISBN 978-9-6331-3140-4. doi: 10.1109/MTITS.2015.7223258. URL <http://ieeexplore.ieee.org/document/7223258/>.
- [55] Michal Maciejewski, Josep Maria Salanova, Joschka Bischoff, and Miquel Estrada. Large-scale microscopic simulation of taxi services. Berlin and Barcelona case studies. *Journal of Ambient Intelligence and Humanized Computing*, 7(3):385–393, jun 2016. ISSN 1868-5137. doi: 10.1007/s12652-016-0366-3. URL <http://link.springer.com/10.1007/s12652-016-0366-3>.
- [56] D. G. Champernowne. A Comparison of Measures of Inequality of Income Distribution. *The Economic Journal*, 84(336):787, dec 1974. ISSN 00130133. doi: 10.2307/2230566. URL <https://academic.oup.com/ej/article/84/336/787-816/5236443>.