Equity Effects of Rare Events on Transportation Network Company and Transit Riders

Destenie Nock[1, 2,] ORCID 0000-0003-1739-7027

Corey Harper[1], ORCID 0000-0003-1956-5258

Jeremy Michalek[1, 2, 2], ORCID 0000-0001-7678-8197

Anna Cobb[3]

Carlos Mateo Samudio Lezcano[1] , ORCID 0000-0002-6797-6329

**FINAL RESEARCH REPORT**

**Grant Number <u>69A3551747111</u>**

Disclaimer: The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated under the sponsorship of the U.S. Department of Transportation's University Transportation Centers Program, in the interest of information exchange. The U.S. Government assumes no liability for the contents or use thereof.

[1] Department of Civil and Environmental Engineering, Carnegie Mellon University
[2] Department of Mechanical Engineering, Carnegie Mellon University
[3] Department of Engineering and Public Policy, Carnegie Mellon University

# Table of Contents

## Overall Description of Project

Increasing mobility for all requires equitable transportation access regardless of location. Within households that do not own personal vehicles, public transit has been instrumental in providing mobility to jobs and other essential services. Now ride-hailing services from transportation network companies (TNCs), like Uber and Lyft, have revolutionized urban transportation by providing an on-demand service option for public transit dependent populations. Yet, the benefits and costs of these changes have been inequitably distributed, widening the gap between those with and without high levels of mobility. This inequity, partially seen in disadvantaged neighborhoods, stems from private ride-hailing firms having profit driven incentives, and less regulation than public transit agencies that need to balance economic and equity objectives in their decision-making.

Lack of reliable and efficient transportation is often cited as a pivotal barrier to healthcare, employment access, and upward socioeconomic mobility. Many people without regular access to automobiles depend on public transit as their main mode of transportation. In densely populated neighborhoods, a fixed route system may work well, since walking distance to bus and train stops may be acceptable. But in medium or low-density areas where residents may have to travel longer distances to and from transit stops, the lack of accessibility creates tremendous challenges for human mobility and leads to usage of unsustainable transportation modes. Recently, transportation network companies (TNCs) (e.g., Uber and Lyft) have revolutionized mobility in many areas by detaching car access from car ownership, and in theory reducing many mobility gaps that arise from people not having access to a personal vehicle.

However, there are economic barriers to these services due to high costs, especially during peak hour travel times, and lack of supply in low-income neighborhoods, resulting in longer wait times and higher cancellation rates. Our investigation identifies opportunities for public policies that may enhance transportation benefits, while mitigating private costs, social costs, and inequities in disadvantaged neighborhoods. To do this, we (1) leverage historical data to econometrically estimate the causal impact of extreme weather events on ride-hailing service operation disruptions and how this effect was distributed across riders served; (2) characterize how the level of service varies by demographic community area; and (3) use simulation and optimization models to identify transportation disparities.

The remainder of the report is structured in two parts. First, we discuss the simulation of racial discrimination in TNCs and the impact this could have on wait times in various communities. Second, we present our weather analysis, and discuss the implications of rare weather events on TNC usage in different communities.

# The Effects of Discrimination in Ridesourcing Systems

Anna Cobb

Department of Engineering and Public Policy,

Carnegie Mellon University, Pittsburgh, PA 15213

Corey Harper, Ph.D. - 0000-0003-1956-5258

Department of Civil and Environmental Engineering,

Carnegie Mellon University, Pittsburgh, PA 15213

Jeremy Michalek, Ph.D. - 0000-0001-7678-8197

Department of Engineering and Public Policy & Department of Mechanical Engineering,

Carnegie Mellon University, Pittsburgh, PA 15213

Aniruddh Mohan, Ph.D. - 0000-0002-6797-6329

Andlinger Center for Energy and the Environment,

Princeton University, Princeton, NJ 08544

Destenie Nock, Ph.D. - 0000-0003-1739-7027

Department of Civil and Environmental Engineering & Department of Engineering and Public Policy,

Carnegie Mellon University, Pittsburgh, PA 15213

## Introduction and Motivation

Taxis have long been known for having discrimination drivers. From the 1980s up to 2016, racial audit studies have revealed that black riders wait 27-52% longer to get picked up by taxis than their white counterparts [1], [2], [3]. However, ridesourcing has evolved substantially in the last decade, as transportation network companies (TNCs) such as Uber and Lyft emerged in cities across the U.S. and around the world. TNCs connect riders to contracted drivers via an app, which allows for features like automatic payment via credit card and star ratings that may be left by both passengers and drivers for their counterparts. Together, these innovations are seen as having reduced some of the motivations for statistical discrimination that may have been present in taxi operations, primarily by reducing the amount of information a driver will try to infer about a potential passenger solely based on their appearance.

In determining the degree to which discrimination may still exist in TNCs despite their operational differences from taxis, two bodies of work have emerged: 1) studies using estimated wait time data from the Uber and Lyft APIs to measure spatial accessibility of TNCs, and 2) TNC-specific audit studies, which attempt to reveal whether drivers practice racially discriminatory behavior. The research discussed in this report builds off results from the second group, which have shown that there is still evidence of discrimination in TNC transactions—albeit substantially less than in taxis. This discrimination largely shows up in the form of cancellations; drivers have been found to be up to twice as likely to cancel rides for black riders than for white riders [2], [3], [4]. Interestingly, the increased cancellation rates for these customers have not been found to result in significantly higher overall wait times. Along with not being able to fully explain this discrepancy, the audit studies are limited by their scope in terms of the number of rides and the types of discrimination they are able to analyze. On the latter point, the scope of audit studies is limited to evaluating direct discrimination by drivers against customers by controlling for ride time and location; in doing this, they are unable to measure differences in wait times that occur because of the correlation between race and where people live in a city. The spatial accessibility studies, on the other hand, are limited by their data; rather than using *experienced* wait times in their analysis, they rely on the *expected* wait times generated by the TNCs' application program interfaces (APIs). While these APIs represent one of the only sources for wait time data, even Uber's former CEO has admitted that Uber's API was "different than actual (i.e. wrong)" [5].

To avoid the limitations imposed by the scope, sample size, and data of previous methodologies used to study this topic, we use an agent-based model (ABM) to simulate millions of recorded TNC trips which occurred in the city of Chicago after 2018. The ABM allows us to simulate these trips under a wide variety of conditions, such as with and without discriminatory drivers, with varying TNC fleet sizes, and with variations in customer-driver matching algorithm priorities. Furthermore, it enables us to investigate at a fundamental level what differences in TNC and taxi company operational structure have dampened the impact of discriminatory drivers on customer wait times in TNCs, as well as how different racial groups are individually impacted by discrimination. Lastly, we use the ABM specifically to identify differences in wait times across racial groups without any discriminatory drivers in the simulation.

## Methods

*Model Overview:*

An ABM written in Julia—a dynamic, high performance computing language—is used to simulate a day's worth of TNC trips in Chicago [6]. Known as AgentX, the model was first created and used by Mohan et al. to compare the externalities of operating a TNC fleet of battery-powered vehicles and gas-powered [7]. The model takes as input a day's worth of recorded Chicago TNC trips from the Chicago TNC Data Portal and an OpenStreetMap file which describes the road network of the city. As demonstrated by its eight second time step, the model is unique in its ability to simulate the movements of thousands of agents at a time at a relatively low computational cost and at high levels of geospatial and temporal detail. Trips are simulated using two types of agents: customers and drivers.

*Customer Agents:*

Each customer agent represents one trip from the day of trips used as input to the model. Customer agent behavior is relatively simple; customers appear in the simulation at their appointed trip start time, wait until a driver arrives to pick them up, and become inactive in the simulation once their destination is reached.

Each customer agent is assigned a race so that the differences in travel experience for riders of different races can be evaluated. As the demographics of TNC riders are not recorded in the Chicago TNC trip data, race is assigned to travelers based on data from the National Household Travel Survey (NHTS) and the Chicago Metropolitan Agency for Planning (CMAP) community data snapshots. As shown in Figure 1, the NHTS data is filtered for respondents from Chicago and processed to reveal the proportion of trips that start at home, end at home, and neither start nor end at home for each hour of the day. Each trip is then assigned one of these three "home statuses" by drawing from the distribution corresponding to the hour of the day in which the trip began. Depending on the home status of the trip, the community area in which the trip originated or ended in is labeled as the home of the rider. Finally, the CMAP data is used

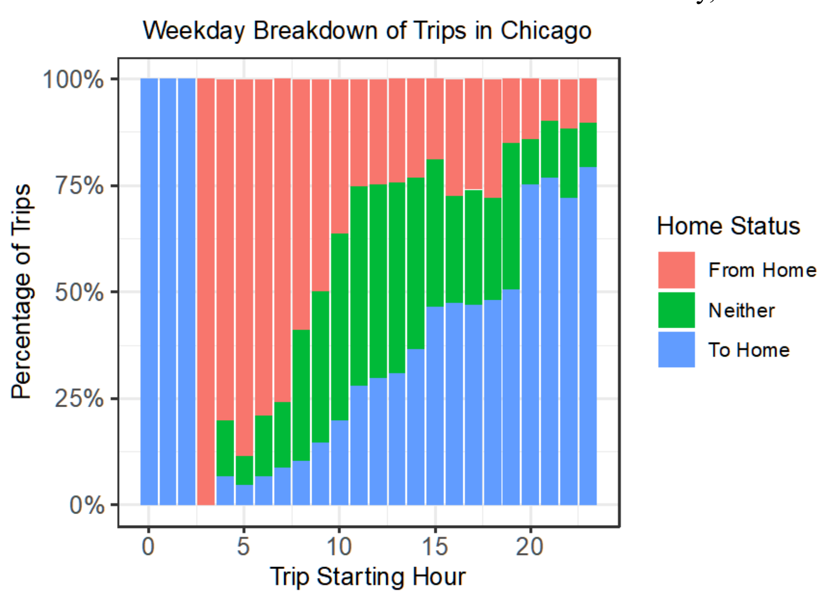


**Figure 1: Distribution of Chicago trips recorded in the 2020 NHTS data which were recorded as starting at home, ending at home, or neither.**

to draw a race from a discrete distribution representing the demographic breakdown of the identified home community area. If a trip is designated as neither starting nor ending at home, or if the trip began in the O'Hare community area (where the O'Hare Airport is located) a race is assigned using the demographic breakdown for the entire city of Chicago.

*Driver Agents:*

Unlike customer agents, the behavior of driver agents is not strictly dictated by empirical data and instead is largely governed by the set of model heuristics described in detail in the subsequent *Model Heuristics and Calibration* section. This is a result of Chicago keeping records of TNC drivers registered to drive in the city, but only at the monthly aggregated level–e.g., instead of providing details on all the trips a given registered driver is responsible for, a monthly count is provided for each registered driver. To assess sensitivity of findings to our modeling heuristics and assumptions, extensive sensitivity analysis will be conducted on each of these heuristics.

Separately from the heuristics, driver behavior can be described at a high level as shown in Figure 2. When a threshold for unmet demand is surpassed within the simulation, drivers appear in the model at a random intersection within city limits. After their initial appearance, drivers either match with a rider and drive towards the rider's pickup location or, if unmatched, drive towards a designated repositioning location. The ABM does not account for traffic flow; each driver agent moves at a speed of 22 mph, roughly the average speed at which TNC trips are completed in Chicago [8]. Drivers exit the simulation when they have met their daily earnings target or have experienced a prolonged period of low revenue. They stop for refueling when fuel reaches below 10% of maximum capacity, which is represented, for simplicity, by ten minutes of no movement in the geographic position in which they dropped below the refueling threshold.

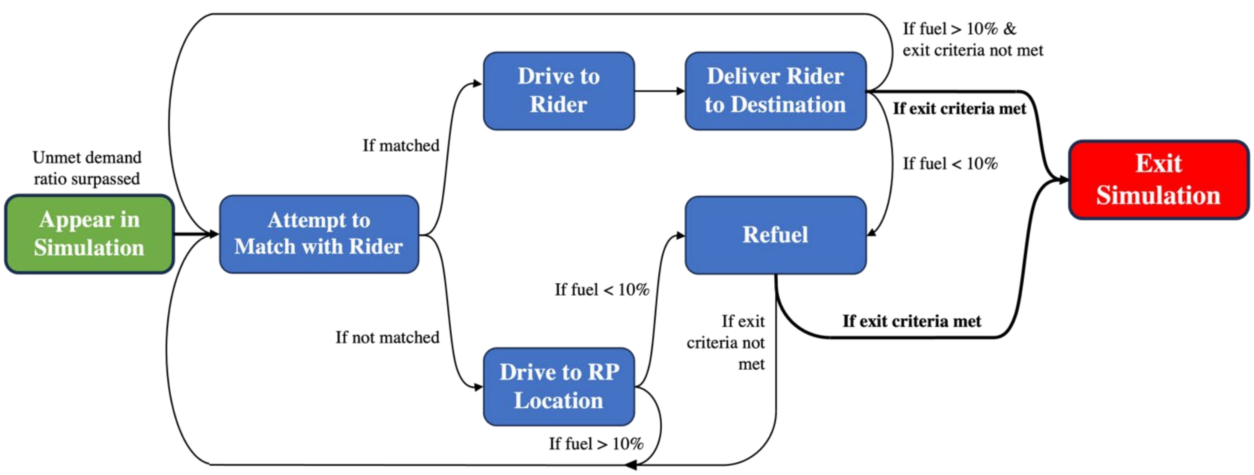


**Figure 2: Flowchart of driver agent behavior during each time step, where "RP location" refers to where a driver is repositioning to and exit criteria are 1) experiencing a prolonged period of low revenue or 2) reaching the driver's earnings target.**

*Model Heuristics and Calibration:*

Four model heuristics govern driver agent behavior and were used to calibrate the model: 1) bringing drivers online, 2) taking drivers offline, 3) fleet-wide driver repositioning strategy, and 4) matching algorithm.

1. **Bringing drivers online:** As previously mentioned, the time at which drivers “come online”, or are added to the simulation, is decided by the current level of unmet demand. Every two minutes, or 15 time steps, the ratio of unmatched customers to available drivers is computed. If this ratio is less than a prescribed limit—1.5 for low trip-count days and up to 3.5 for high trip-count days—additional drivers are added to the simulation until the ratio drops below the limit. These limits were calibrated to ensure that 95% of customers were served in 20 minutes or less, a target that is derived from both TNC modeling literature and anecdotal evidence of wait times being longer than 5-10 minutes [9], [10].

2. **Drivers going offline:** Drivers exit the simulation, or “go offline”, for two potential reasons: they have reached their daily earnings target or they have experienced a prolonged period of low earnings. Drivers’ earnings targets are drawn from a discrete distribution of earnings values computed such that aggregate driver daily working hours correspond with a pre-defined distribution: 50% of drivers work less than four hours, 30% work four to six hours, and 20% work seven or more hours. This distribution of daily working hours was created with the intention of combining results from a 2020 report by Uber which stated that only 9% of drivers in California spend 40 or more hours a week on the app [11] and a 2020 study set in Seattle which found that one third of Uber drivers work more than 32 hours weekly [12]. Typically, less than 10% of drivers go offline as a result of the low-revenue period criteria, enabling calibration of driver working hours largely by modifying the earnings targets assigned to drivers.

3. **Driver repositioning strategy:** We have altered driver repositioning strategy from its original implementation by Mohan et al. [7] to be more localized and thus, more representative of experienced drivers. To accomplish this, Chicago TNC trip data for the month and year corresponding to each day of trips that is simulated is processed to yield the five community areas with the highest counts of trip pickups for each hour of the day. The results are used as an additional input to the simulation, and for a given hour, driver agents pick the closest community area of the top five to drive towards when repositioning.

4. **Matching algorithm:** Drivers and customers are matched with one another first using a 5 km radius, then expanding to a 10 km radius if no matches are found. If either rider or driver is matched to multiple other drivers or riders, respectively, the distance between rider and driver, as well as customer wait time and driver revenue, are used to make a singular match. Though TNC matching algorithms remain oblique to the public, equalizing revenue across drivers and minimizing customer wait times is in line with findings from the literature [13], [14] and with what is known about Uber’s "Batch Matching" strategy [15]. Our implementation of driver discrimination against customers comes into play during the matching process. We designate discriminatory drivers by adding a dummy "bias" attribute to each driver agent. When a driver’s bias attribute is set to 1, they will reject their customer match if they are biased against the customer’s race. The customer’s ID is recorded such that the same customer will not be suggested to the driver again and the driver is suggested a new customer in the next time step.

## Data

It should be noted that while Chicago was originally chosen because of the detailed data the city publishes on TNC trips and drivers, it also happens to have substantial amounts of residential segregation. For community areas, the geographic region which is used by the city government for data collection and statistics, in which Black residents make up more than 50% of the population, they make up 86% of the community area on average; when less than 50%, Black residents make up an average of 10%. This makes Chicago an especially well-suited city for our analysis of travel differences related to race and geography.

***Chicago TNC Trip Data*** [16]

The trip data used as input to the simulation comes from the Chicago Data Portal, which has maintained records of every TNC trip that starts or ends within city limits beginning in November 2018. Trip start and end time are reported in 15 minute blocks and trip start and end location are recorded at at least the community area level. Additional information, such as the distance, travel time, and fare of the trip are also recorded. For each day which is simulated, 26 hours' worth of trip data are downloaded, beginning at 2:30 AM on the day of interest and ending at 4:30 AM on the next day. Though the full 26 hours of trips are simulated, only the middle 24 hours (3:30 AM to 3:30 AM) of trips are post-processed. The first hour of trips (2:30-3:30 AM) are not post-processed to avoid including trip data from when the model is still reaching its steady state. The last hour (3:30-4:30 AM) is only simulated to ensure that all trips started during the last hour of the desired 24-hour window (2:30-3:30 AM) are completed, and so data on the trips beginning during that hour do not need to be post-processed.

*Chicago TNC Trip Data Preprocessing:*

The origins and destinations of trips that begin or end outside of city limits are not recorded, and therefore are removed from our 26 hours of input trips. This process reduces total trip count by about 10%. Next, each trip is randomly assigned a start time within its listed 15-minute block and a starting location at an intersection within its listed origin community area. Both random assignments are done using uniform distributions. Finally, each trip, which will eventually be representative of a customer agent, is assigned a race.

***National Household Travel Survey (NHTS)*** [17]

As discussed in the Customer Agents section, the NHTS data is filtered for Chicago respondents, then used to classify each input trips as likely to have started at home, ended at home, or neither, dependent on the hour in which the trip began. This means the demographic breakdown of the origin or destination community area (or the city of Chicago if classified as neither) can be used to assign race.

***CMAP Community Data Snapshots*** [18]

As discussed in the Customer Agents section, demographic data from the CMAP community area snapshots is used to assign a race to each trip in our input set based on home status assignment using the NHTS data.

## Results (Preliminary)

*Proportion of Biased Drivers*

We began by investigating what percentage of drivers would need to be biased to observe the cancellation rates found in the audit studies. Specifically, we focus on the audit conducted in Boston by Ge et al. (2), which found a difference of 5.1% between the average cancellation rate of Black (10.1%) and white (4.9%) Uber riders. Since no drivers cancel on riders during the simulation unless racially biased against them, the percentage of biased drivers can be incrementally increased until the cancellation rate against Black customer agents is measured to be 5.1% on average. We found that this occurs at around 3.3% of driver agents being biased, though the percentage varies slightly across the dates being simulated. As it is certainly possible that rather than there being a small percentage of consistently biased drivers, there is instead a larger percentage of drivers who inconsistently exhibit bias, this value provides a lower bound for what proportion of TNC drivers at least sometimes cancel on riders because of their race.

*Wait Times across Racial Groups*

Using a value of 3.3% for the proportion of biased drivers, Figure 3 shows how wait times are impacted across racial groups when bias is present in the simulation and when it is not for October 10th, 2021.

In addition to running a scenario in which 3.3% of drivers were only biased against Black customers, a scenario was run in which the same percentage of drivers were biased against two racial groups: Black and Hispanic customers. Though the only audit study that checked for discrimination

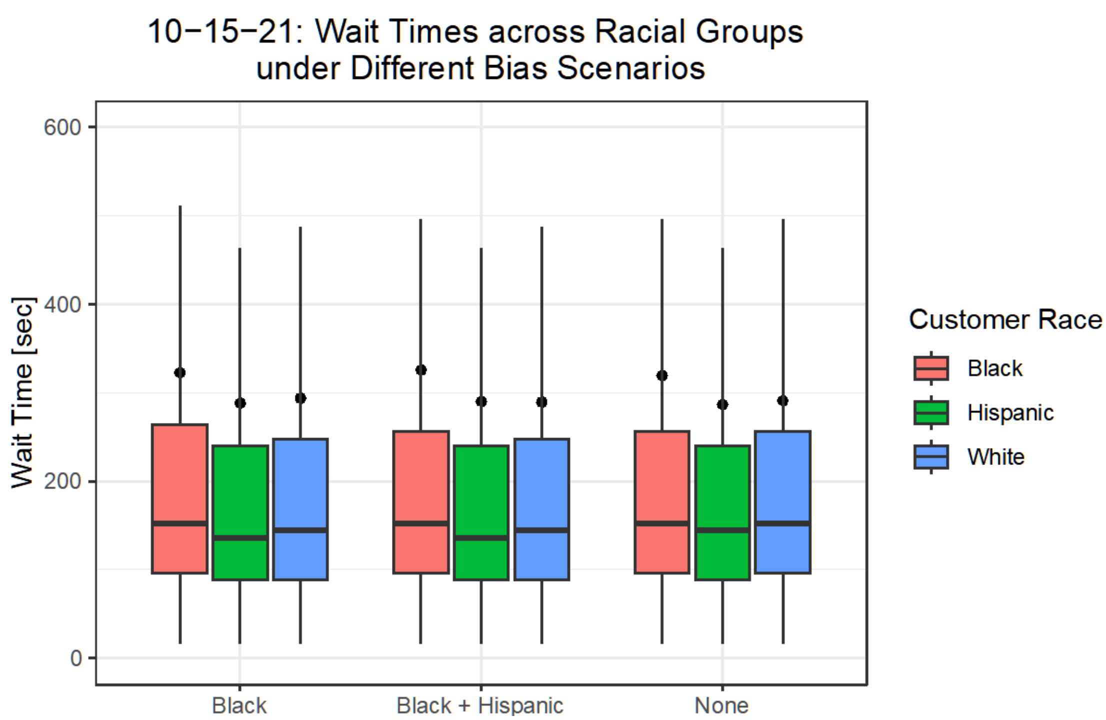


**Figure 3: Distributions of wait times across racial groups under three bias scenarios: 3.3% of drivers biased against Black customers, 3.3% of drivers biased against Black and Hispanic customers, and no bias present. Means are shown as circles and outliers are not plotted.**

against Hispanic riders did not find the same effect in Uber cancellation rates as with Black riders [3], this is not sufficient evidence to rule out the idea that drivers who discriminate against Black riders would not potentially discriminate against other people of color. However, it should be noted that after running this scenario, the cancellation rates measured for Black customers increased substantially, up to 10%. We hypothesize that this is a result of biased drivers having increasingly limited options for who they are willing to pick up and thus where they are willing to drive, meaning they more often get stuck in neighborhoods where they refuse to pick customers up.

From Figure 3, it is clear that when 3.3% of driver agents practice discriminatory behavior, no racial group's average wait time is substantially impacted. In fact, the median wait times for each racial group stay the same across all bias scenarios with the exception of white customer wait times in the Black and Hispanic bias scenario, which decreases by eight seconds. As shown in Table 1, average wait times capture small changes in outlier behavior; for October 10th, 2021, as well as each of the other dates simulated, there is an increase of two to three seconds in the average wait time of Black customers when drivers are biased against only them. This very small but consistent increase in wait times is in line with the audit studies discussed earlier. We hypothesize that this is because of the oversupply of drivers in the simulation and relatively quick matching and re-matching time (every eight seconds). As a check on the former point, simulations were run for October 10th, 2021 in which the unmet demand threshold used to bring taxis online was increased in increments of 0.1 from 1.6 to 2.3 (or 16 to 23 available taxis per 10 unmatched customers). As the relative oversupply of taxis was increased, the percent difference in average wait times of Black and white customers with 3.3% discriminatory drivers in the simulation decreased from 9% to 6%. This helps to explains why the effects of discrimination by taxi drivers were felt so much stronger in audit studies, as taxis are usually largely undersupplied in cities.

**Table 1: Average wait times for white and Black customers for simulations with no biased drivers.**

| Bias Scenario | | | Average Wait Time for Racial Groups [sec] | | |
|---|---|---|---|---|---|
| **Date** | **Bias Against** | **Biased Drivers (%)** | **White** | **Black** | **% Difference (Black vs. White)** |
| 2/14/23 | Black + Hispanic | 3.33 | 351.03 | 374.13 | 6% |
| | Black | 3.33 | 352.07 | 374.51 | 6% |
| | N/A | 0.00 | 350.92 | 372.53 | 6% |
| 10/15/21 | Black + Hispanic | 3.33 | 289.34 | 325.47 | 11% |
| | Black | 3.33 | 293.75 | 322.45 | 9% |
| | N/A | 0.00 | 291.04 | 319.25 | 9% |
| 7/26/21 | Black + Hispanic | 3.33 | 354.61 | 395.21 | 10% |
| | Black | 3.33 | 359.23 | 396.95 | 10% |
| | N/A | 0.00 | 363.75 | 396.58 | 8% |

Secondly, there is a clear trend in the ordering of wait times across racial groups for each date, regardless of whether or not there are discriminatory drivers present. While the differences in wait times resulting from the addition of discriminatory drivers to the simulation were only a few seconds, average wait times for Black customers are 21.61 - 32.83 seconds longer than for white riders when no bias is present for the dates shown in Table 1. To understand if this trend could be consistently observed, a larger array of dates were simulated without any discriminatory drivers in the simulation. Results from these additional days are shown in Table 2. With the exception of July 14th, 2022, the positive difference in average wait times for Black and white customers indicate that there is a connection between the geographic location in which Black and white riders request rides and the availability of TNC drivers. This finding is in contrast to the conclusions of TNC spatial accessibility studies, which did not find a link between a region's racial composition and its expected TNC wait times.

## Future Work

There are several items which will be investigated as a part of the future work to be done for this project.

1. We will incorporate Chicago's TNC Driver data into our simulation. Specifically, this will be done by processing registered drivers' residential zip codes (included in the dataset) so that they may be used to inform where drivers appear in the simulation. Additionally, the driver dataset—updated monthly—includes the number of trips each driver completed in a given month, which could be used to improve our distribution of daily hours worked by drivers.

2. We will run a number of sensitivity cases on our model heuristics to understand under what circumstances a TNC operating structure does show sensitivity to discrimination. As mentioned in the introduction, this will include increasing the time step of the ABM to slow down the rate at which drivers are matched to new customers, varying the weights on customer wait time and driver revenue in the matching, implementing additional repositioning strategies, and potentially decreasing the unmet demand limit at which we add more drivers to the simulation. Through this process, we aim to better which features of a TNC operating system are impact wait time sensitivity to discriminatory drivers.

3. As an additional sensitivity check, we will run all simulation days with trips starting and ending at the O'Hare International Airport removed. There are two motivations for this, the first being that those taking Ubers to and especially from the airport may not be representative of the general Chicago population. Chicago is a business, academic, and tourism hub, and it is likely that many of the people taking Uber rides from the airport are not Chicago residents (making our assumptions about their race invalid). Secondly, the O'Hare Airport represents an anomaly, in that it is roughly 17 miles from trip-dense downtown Chicago, but 5-10% of TNC trips in Chicago either originate or end there. This means that it is hard to simulate real-life driver behavior in which drivers make the long, unprofitable trip out to the airport and wait however long is required to pick up a customer. It should be noted that the same issue does not apply for Midway International Airport, which is significantly smaller and located closer to the city center.

## Conclusions

As TNCs continue to dominate taxis in the ridehailing market, it is important to understand how the effects of discriminatory practices by drivers against customers manifest themselves in this new type of operating structure. In this work, we have filled a gap in the existing literature by using an agent-based model to simulate TNC trips in Chicago, allowing us to expand the scopes of methodologies previously used to study discrimination in TNCs. Our preliminary results indicate that the biggest driver in wait time gaps between white and Black TNC customers—up to a 15% increase for Black customers—in Chicago come from residential segregation in the city, rather than from discriminatory drivers. This assumes that the level of racial discrimination by TNC drivers against customers in the form of increased cancellation rates is not significantly higher in Chicago than in the three cities where TNC racial audits have been performed. While the small increase in wait times observed as a result of driver discrimination is in line with findings from audit studies, the more systemic differences in wait times observed without bias present in the system do not align with the results of spatial accessibility studies. Meaningful differences between wait times of racial groups would be a strong reason to begin regulating TNCs in a more similar fashion to taxis; whereas many taxi companies are subject to regulations requiring them to serve all areas of a jurisdiction regardless of where demand is highest, TNCs are not currently subject to such rules (35).

However, it should once again be noted that all results shown in this report are preliminary and should be treated as such. Additional sensitivity analysis will be conducted to ensure that model heuristics are not driving the differences we have observed in wait times and where possible, additional data will be incorporated into the model to ensure that reality is reflected as accurately as possible.

## Acknowledgements

The authors would like to thank Professor Don Mackenzie at the University of Washington and Professor Anne Brown at the University of Oregon for discussing their audit study papers with us. This research was also supported by the US DOT [Grant Number 69A3551747111] through Mobility21, a University Transportation Center, with the goal of improving mobility of goods and services.

## References

[1] P. Siegelman, “Racial discrimination in ‘everyday’ commercial transactions: What do we know, what do we need to know, and how can we find out,” A national report card on discrimination in America: The role of testing, pp. 69–98, 1998.

[2] Y. Ge, C. R. Knittel, D. MacKenzie, and S. Zoepf, “Racial discrimination in transportation network companies,” Journal of Public Economics, vol. 190, p. 104205, 2020, doi: https://doi.org/10.1016/j.jpubeco.2020.104205.

[3] A. E. Brown, “Prevalence and Mechanisms of Discrimination: Evidence from the Ride-Hail and Taxi Industries,” Journal of Planning Education and Research, vol. 43, no. 2, pp. 268–280, 2023, doi: 10.1177/0739456X19871687.

[4] Jorge Mejia, Chris Parker (2021) When Transparency Fails: Bias and Financial Incentives in Ridesharing Platforms. Management Science 67(1):166-184. https://doi.org/10.1287/mnsc.2019.3525

[5] “Yes, Your Uber Wait Times Are Almost Always Wrong. Here’s Why.,” *Thrillist*. https://www.thrillist.com/tech/nation/why-uber-wait-times-are-always-wrong (accessed Jun. 07, 2023).

[6] Datseris, G., A. R. Vahdati, and T. C. DuBois, Agents.jl: A performant and feature-full agent based modelling software of minimal code complexity, 2021.

[7] Mohan, A., M. Bruchon, J. Michalek, and P. Vaishnav, Life Cycle Air Pollution, “Lifecycle Air Pollution, Greenhouse Gas, and Traffic Externality Benefits and Costs of Electrifying Uber and Lyft, “Environmental Science & Technology, Vol. 57, No. 23, 2023, pp. 8524–8535, pMID: 37260172.

[8] Schneider, T.W., “Taxi and Ridehailing Use in Chicago,” https://toddwschneider.com/dashboards/chicago-taxi-ridehailing-data/, 2023 (accessed: May 05, 2023).

[9] Boesch, P. M., F. Ciari, and K. W. Axhausen, “Autonomous Vehicle Fleet Sizes Required to Serve Different Levels of Demand. Transportation Research Record,” Vol. 2542, No. 1, 2016, pp. 111–119.

[10] Allyn, B., “Lyft And Uber Prices Are High. Wait Times Are Long And Drivers Are Scarce,” NPR, 2021.

[11] Mishkin, L., “Which drivers do the most trips?” Uber, 2020.

[12] Parrot, J. A. and M. Reich, “A Minimum Compensation Standard for Seattle TNC Drivers,” *Center for New York City Affairs*, 2020.

[13] Hall, J. V. and A. B. Krueger, “An Analysis of the Labor Market for Uber’s Driver-Partners in the United States,” ILR Review, Vol. 71, No. 3, 2018, pp. 705–732.

[14] Chen, M. K., J. A. Chevalier, P. E. Rossi, and E. Oehlsen, “The Value of Flexible Work: Evidence from Uber Drivers. Journal of Political Economy,” Vol. 127, No. 6, 2019, pp. 2735–2794.

[15] Uber, "How does Uber match riders with drivers?"
https://www.uber.com/us/en/marketplace/matching/, 2023 (accessed: Jul. 09, 2023).

[16] "Transportation Network Providers – Trips (2018-2022)," Chicago Data Portal. 2022 [Online]. Available: https://data.cityofchicago.org/Transportation/Transportation-Network-Providers-Trips-2018-2022-/m6dm-c72p

[17] "2017 National Household Travel Survey," U.S. Department of Transportation. 2017 [Online]. Available: https://nhts.ornl.gov.

[18] "Community Data Snapshots," Chicago Metropolitan Agency for Planning. July 2022 [Online]. Available: https://www.cmap.illinois.gov/data/community-snapshots

# Equity Under the Effects of Weather on Transportation Network Companies

Carlos Mateo Samudio Lezcano

Department of Civil and Environmental Engineering,

Carnegie Mellon University, Pittsburgh, PA 15213

Corey Harper, Ph.D. - 0000-0003-1956-5258

Department of Civil and Environmental Engineering,

Carnegie Mellon University, Pittsburgh, PA 15213

Destenie Nock, Ph.D. - 0000-0003-1739-7027

Department of Civil and Environmental Engineering & Department of Engineering and Public Policy,

Carnegie Mellon University, Pittsburgh, PA 15213

Jeremy Michalek, Ph.D. - 0000-0001-7678-8197

Department of Engineering and Public Policy & Department of Mechanical Engineering,

Carnegie Mellon University, Pittsburgh, PA 15213

## Introduction

The use of Transportation Network Companies (TNCs) has become an essential part of urban transportation in recent years. However, very little is known on the vulnerability of prices to weather events like heavy rain and snowfall. Our work leverages the vast amounts of publicly available TNC trip data and weather data to explore the effects of rain on trip price. This builds on previous studies about the equity in TNC operations [4,5], and impacts of external shocks on ridership in shared transportation methods [6].

## Data

The weather data used for this work consists of hourly records of rain intensity, snow intensity and apparent temperature. This data was gathered from the DarkSky API (*Darksky*, 2021), which provides historical weather data. As for the TNC trip data, we the Transportation Network Providers Trip dataset (Chicago Data Portal, 2023). This dataset contains all trips made by Transportation Network Companies (TNCs) in Chicago from the period of January 2019 to December 2019.

## Case Study

This study takes place in the city of Chicago. Chicago is located in the state of Illinois, at its eastern border with the state of Michigan. Chicago is also located next to Lake Michigan. According to the U.S. Census Bureau, the city has an area of 12,059 square miles and a population of 2,665,039 residents. The median household income is 65,781 $, with 41,7% of residents 25 and above holding a higher education degree. The city of Chicago's transit system consists of an extensive network of buses, trains and subways.

## Methodology

In order to estimate the effects of different weather events on price and ridership of TNCs, we estimate the Average Treatment Effect of rain on trip prices, using the Augmented Inverse Propensity Weighted estimator, as defined in Equation 1.

*Equation 1. Average Treatment Effect*

$$\widehat{\mathrm{ATE}} = \frac{1}{n}\sum_{i=1}^{n}\left(\widehat{\mu_{(1)}}(X_i) - \widehat{\mu_{(0)}}(X_i) + W_i\frac{Y_i - \widehat{\mu_{(1)}}(X_i)}{\hat{e}(X_i)} - (1 - W_i)\frac{Y_i - \widehat{\mu_{(0)}}(X_i)}{1 - \hat{e}(X_i)}\right)$$

Detailed treatments of this estimate and its properties can be found elsewhere in the literature (Kennedy, 2022). Briefly, n is the number of samples, $X_i$ is the vector of covariates for datapoint i, $W_i$ is a binary variable equal to 1 when datapoint i received treatment, and 0 otherwise, $Y_i$ is the observed outcome, $\hat{e}(X_i)$ is the estimator of the propensity score (i.e., the probability of treatment conditional on the observed covariates), and $\mu_{(j)}(X_i)$ is the estimator of the potential outcome conditional on the observed covariates, with j = 1 corresponding to the potential outcome if treated, and j=0 corresponding to the potential outcome if untreated.

## Potential Outcome Estimators Performance

To estimate the potential outcome models of trip price during different weather events, we fit different regression models to the individual trip data. One model ($\mu_{(0)}$) is fitted on a control dataset, where the datapoints did not receive treatment, and another model ($\mu_{(1)}$) is fitted on a treatment dataset, where the datapoints received treatment. Table 1 shows the definition and characteristics of the treatment and control groups, for the different weather events. The resulting metrics for the different potential outcome and propensity score regression models are shown in Table 2, with the bolded entries corresponding to the lowest achieved mean squared error (MSE).

**Table 1. Characteristics of treatment and control groups for price models**

| | Rain | |
|---|---|---|
| | **Treatment** | **Control** |
| **Number of Samples** | 1,161,568 | 77,265,653 |
| **Precipitation Intensity** | > 2.5 mm | 0 |

**Table 2. Test loss metrics for price regression models. Note: the propensity score estimator was a logistic regression model, so we used accuracy as the loss. The potential outcome models use MSE.**

| | Rain | | |
|---|---|---|---|
| | **μ0** | **μ1** | **e** |
| **Logistic Regression** | - | - | 0.974 |
| **Boosted Decision Trees** | 0.051 | 0.057 | - |

## Results (Preliminary)

The results of the estimation of average treatment effects of rain on price conditioned on trip origin Community Area are shown in Figure 1. The point estimates range from decreases of 2.51% to increases of 21.11%. The results show an interesting spatial pattern, where the largest price increases are clustered in the areas directly north of Downtown, which coincides with the areas with the highest percent of employed residents (**Error! Reference source not found.**, part (a)). These same areas experienced very little effect on ridership due to rain, so trips originating in these areas are less sensitive to the effects of rain, suggesting that many of these trips are work related or commute trips. The effects on price tend to attenuate, and then change sign to become price decreases as we move away from the Downtown area. The change in effect is gradual, and we can also observe that areas near the lake front experience price increases more so than the others.

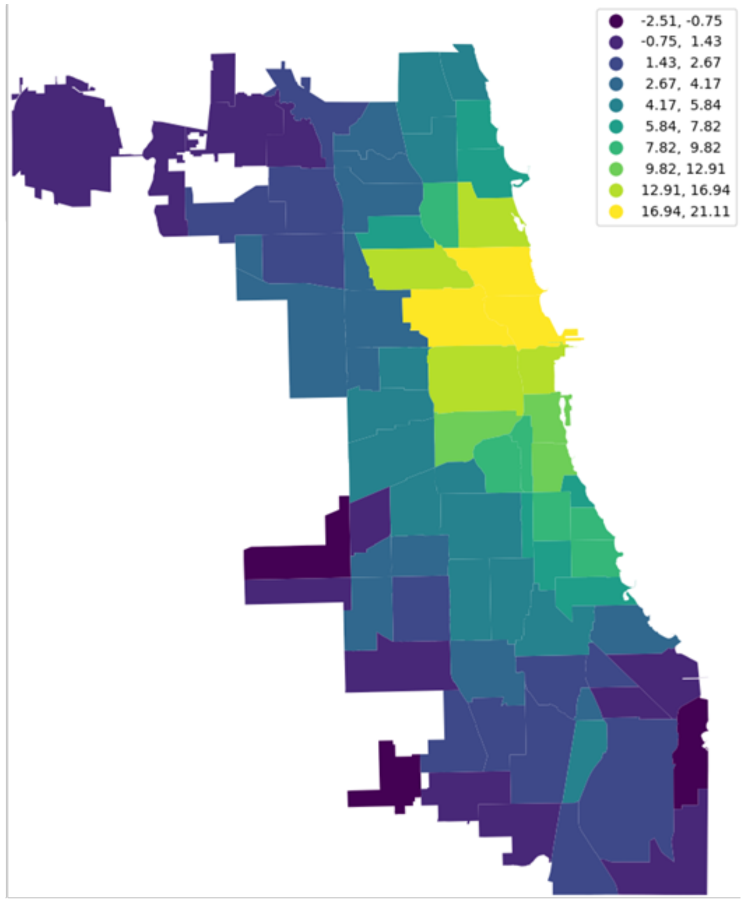


(b)

**Figure 1. Average Treatment Effect of Rain on Trip Price in % Change, Conditioned on Origin of the Trip. Period: 7-10 am.**

## Conclusions

The goal of this research was to explore the effects of different weather events on TNC ridership and prices by leveraging the vast amounts of existing data, so as to provide useful insights for transportation planning practitioners. We have achieved this by combining TNC trip data, weather, demographic, and travel pattern data, and building models with all this data to estimate the average treatment effects of these events on TNC ridership and price.

We see that rain a high on prices, whereas snow has the most impact on demand. This indicates that activities are not usually cancelled due to rain in Chicago, so when supply is disrupted due to reluctance to driving in the rain, ridership remains close to its undisrupted levels and hence prices spike at the community area level up to 21.11 %.

Future work will focus on quantifying the effects of rain and other weather events on ridership as well as price. In addition to this, once the effects of weather on price and ridership are estimated, the spatial difference in effects will be analyzed against the spatial demographic distribution, to quantify the difference in effects for different demographics.

## Acknowledgments

This research was supported by the US DOT [Grant Number 69A3551747111] through Mobility21, a University Transportation Center, with the goal of improving mobility of goods and services.

## References

[1] Chicago Data Portal. (2023). *Transportation Network Providers—Trips*. https://data.cityofchicago.org/Transportation/Transportation-Network-Providers-Trips-2018-2022-/m6dm-c72p

[2] *Darksky*. (2021). [Computer software]. Darksky. https://darksky.net/forecast/40.7127,-74.0059/us12/en

[3] Kennedy, E. H. (2022). *Semiparametric doubly robust targeted double machine learning: A review*. https://doi.org/10.48550/ARXIV.2203.06469

[4] Pandey, A. & Caliskan, A. Iterative Effect-Size Bias in Ride-hailing: Measuring Social Bias in Dynamic Pricing of 100 Million Rides. ArXiv (2020).

[5] Chen, L., Mislove, A. & Wilson, C. Peeking Beneath the Hood of Uber. Imc'15: Proceedings of the 2015 Acm Conference on Internet Measurement Conference, 495-508, doi:10.1145/2815675.2815681 (2015).

[6] Hanig, L., Harper, C. D., & Nock, D. (2023). COVID-19 public transit precautions: Trade-offs between risk reduction and costs. Transportation Research Interdisciplinary Perspectives, 18, 100762.