Economic Disparities in Ride-Sharing Pricing

Evidence from a Congestion Tax Reform in Chicago

Zhewei Xie

Submitted to Prof. Cody Tuttle
The Assistant Professor of The University of Texas at Austin

Date: April 27, 2025

Department of Economics
The University of Texas at Austin

Contents

Abstract

The paper uses Chicago's 2020 ride-hailing congestion tax as a quasi-experimental shock and applies a Difference-in-Regression-Discontinuity-in-Time (Diff-in-RDiT) design with a weekday-aligned 2019 placebo window. Trip-level Transportation Network Provider (TNP) data are merged with census-tract socioeconomic indicators, transit and taxi activity, and weather controls. The results show that fare increases were largest in high-income areas and smallest in low-income areas—approximately +\$1.97 (high), +\$1.83 (middle), and +\$1.26 (low) within a ±15-day window under a linear trend. These patterns are robust across bandwidth choices, polynomial trends, and log-price specifications. The distribution of fare changes is consistent with proxy-based third-degree price discrimination, where platforms raise prices more in higher-income neighborhoods.

1 Introduction

There has been a substantial growth in the ride-hailing industry since Uber introduced its services in 2009. Since then, ride-hailing has become increasingly popular worldwide and has changed our daily lives and mobility modes. Now, in 2025, the revenue of the ride-hailing industry is estimated to be about \$175.73 billion worldwide[1], while the traditional taxi market is valued at approximately \$139.59 billion[2]. As a phenomenon involving multiple domains, ride-hailing has attracted attention across different fields, including computer science, operations research, transportation studies, economics, and the social sciences. Within economics, discussions have focused on industry mechanism design, allocation efficiency, and consumer welfare. However, the concern about whether the algorithm is biased and expands the gap of the existing inequality has arisen in recent years. After big data technology and artificial intelligence have made great progress, this topic is more likely to provoke people's sensitive nerves.

Before concerns about the platform's mechanics are raised, it is important to acknowledge that ride-sharing offers an option for travelers and aligns with several previously discussed benefits, such as promoting the utilization of idle private vehicles [Agatz et al., 2012] and serving as a complement to existing travel modes [Hall et al., 2018]. However, after looking deeper into its operation mechanism, the core of ride-hailing services, which lies in matching consumers with suppliers, one of the classic challenges in this on-demand business model, has some possible guilt. Many ride-hailing platforms use a technology known as dynamic pricing algorithms to balance demand and supply [Banerjee et al., 2016]. Although they continue to claim that these algorithms can enhance the welfare of both drivers and riders, while also improving overall system efficiency[3], some critics argue that dynamic pricing may reflect existing inequalities in transportation and even widen the gap already existing in society and may risk undermining the broader pursuit of fairness and justice [Pandey and Caliskan, 2021]. One potential way for platform-based discrimination is price discrimination. Although regulatory restrictions exist, platforms may still implement differential pricing by using proxy information. To examine whether the platform engages in price discrimination against consumers based on pick-up location socioeconomic information, I combined ride-hailing trip data (TNP data) from the City of Chicago with the American Community Survey (ACS), as well as the public transit data, taxi data, and weather data. To groundedly identify potential causal effects, I applied the Difference-in-Regression Discontinuity-in-Time (Diff-in-RDiT) approach to a pooled dataset combining a treated year and a placebo year around a natural experiment stemming from the implementation of a congestion tax on ride-hailing in 2020. Through estimating the effects across heterogeneous groups, I found a systematic pattern: prices increase the most in high-income neighborhoods, while price changes in lower-income areas are minimal, raising concerns about potential algorithmic bias.

To fully understand how such differentiated impacts may arise, it is useful to revisit the underlying business model of the sharing economy and its operational foundations. Introducing these business features provides a natural starting point before studying the core research question. In short, as two-sided markets, where a platform connects and matches demand with supply, sharing services charge consumers on one side while pay fares to service providers on the other [Banerjee et al., 2016]. Such business models have been operating

[1]Ride-hailing Market Insight: https://www.statista.com/outlook/mmo/shared-mobility/ride-hailing/worldwide
[2]Taxi Market Insight: https://www.statista.com/outlook/mmo/shared-mobility/taxi/worldwide
[3]The Effects of Uber's Surge Pricing: A Case Study: https://www.uber.com/blog/research/the-effects-of-ubers-surge-pricing-a-case-study/

for years across various fields, including short-term rentals like Airbnb, daily services like TaskRabbit, and the focus of this paper—ride-hailing services. A key difference between ride-hailing services and earlier forms of sharing services is that ride-hailing is an on-demand product. This means that demand is real-time, and the platform must match supply and demand accordingly. Drivers do not follow a fixed work schedule, and demand itself is influenced by a variety of factors, many of which are unpredictable [Chen, 2016]. These features present significant challenges for platform firms.

To address this challenge, ride-hailing platforms typically use two key technologies: matching and dynamic pricing. In simple words, the former refers to the process of dispatching available drivers to riders, and the latter adjusts ride prices based on demand and supply conditions in real-time. Matching is typically implemented as an algorithm based on space localization and time window limitation, meaning that drivers and riders are matched within a nearby area and a defined time window. Dynamic pricing is implemented as "surge pricing" by Uber and "prime time" by Lyft, which are designed to maximize driver utilization, shorten rider wait times, and ultimately increase the number of successfully matched trip pairs[4], which in turn led to higher profits.

Although these companies have repeatedly emphasized that their algorithms are objective and unbiased[5], there still exists public concern that these "black box" systems may be treating users differently without realizing it. Such concerns have intensified as the public has become increasingly aware that their data may be used to analyze their willingness to pay. One prevalent worry is that these platforms can use travelers' data to construct individual user profiles, which may be applied in business analysis and the implementation of personalized pricing menus [Naini et al., 2016]. These concerns have led to the introduction of regulations aimed at restricting the use of personal information, such as the General Data Protection Regulation (GDPR)[6] in the EU and the American Privacy Rights Act (APRA)[7] in the U.S. However, a profound concern is whether platforms treat different social or demographic groups unequally. Unequal treatment at the group level, such as systematic price discrimination based on income, race, or geographic location, raises concerns about social fairness, structural bias, and discrimination, and has therefore become a deeper concern for the public.

The remainder of this paper is structured as follows. Section 2 reviews a series of studies related to ride-hailing services and concerns about algorithmic bias. Section 3 introduces the classic models of price discrimination, along with their variants that reflect the real-world constraints faced by platforms under regulatory restrictions. Section 4 describes what and how the data I used in this study, including the source of the datasets and the limitations I faced, as well as the manipulation I have done. Section 5 presents the identification framework and outlines the key assumptions that need to be considered and tested. Section 6 reports the empirical results from the natural experiment and provides a discussion of the potential mechanisms underlying these findings.

2 Literature Review

A rich literature motivates my research interest in algorithmic bias in the ride-hailing industry, and concerns about algorithmic bias are not new and have been explored across disciplines. The first thing noticed in prior research was a field study conducted by Sun et al. (2020) in China, revealing device-based disparities in pricing: iPhone users were systematically offered comfort vehicles at a higher price and received more optimistic waiting time estimates. While this study highlights bias at the individual level, it does not investigate whether these disparities can exist at the community level, and this study investigated in China, raising doubts about whether the findings are applicable to the U.S. context. Additionally, they used a field study based on about 800 surveys, which limits the amount of sample in their research. In contrast, my study investigated a similar topic on a combination of millions of records, which can provide more powerful evidence of the widespread existence of algorithm bias. Similar surveys have also been conducted in the U.S. Chang, Winston, and Yan (2022) showed that UberX fares increase with the rating star of hotel destina-

[4]Surge Pricing 101: https://www.uber.com/en-IN/blog/ludhiana/surge-pricing-101/
[5]How surge pricing works: https://www.uber.com/us/en/drive/driver-app/how-surge-works/
[6]General Data Protection Regulation: https://gdpr-info.eu
[7]The American Privacy Rights Act: https://www.congress.gov/crs-product/LSB11161

tions, suggesting price discrimination based on travelers' geographic information [Chang et al., 2022]. My study expands on their research by using daily ride-hailing records to capture price discrimination patterns embedded in everyday urban mobility, while their setting focuses on occasional travel from airports, limiting generalizability. Some researchers do not want to limit themselves to exposing inequality on platforms but hope to gain insights behind the unequal treatment on platforms as the broader social systemic discrimination. Ge et al. (2016) used field experiments to identify differences in ride acceptance rates based on race and gender, reflecting bias at the supply side [Ge et al., 2016]. While valuable, their work did not address how pricing algorithms themselves respond to consumer characteristics. The difference between them is that my study shifts the focus to platform-level decision-making, revealing how pricing adjusts in response to socioeconomic signals, even under a uniform policy. In addition, there is a paper that is most similar to the topic I conducted, contributed by Pandey and Caliskan (2021), who documented group-level price disparities across Chicago neighborhoods [Pandey and Caliskan, 2021]. However, their analysis is correlation-based and does not establish causality. My paper builds on this foundation by introducing a Diff-in-RDiT identification strategy, allowing for solid causal estimation of algorithmic heterogeneous responses on different income groups to exogenous policy shocks. Taken together, this paper contributes to the literature by combining a novel identification strategy with a theoretically grounded framework of proxy-based Bayesian pricing, offering new evidence that algorithmic price discrimination can emerge not from observed preferences, but from community-level proxies inferred by the platform.

In addition to the literature I just mentioned that directly inspired the research interests of this study, there are some other literature that supplemented my understanding of this research field. Let me first go back to the core issue of the online ride-hailing industry mentioned earlier: how to match and how to price. A valuable clue comes from Uber's engineering team. Chiwei Yan et al. (2018) introduced matching solution algorithms such as First-Dispatch Protocol, Maximum Dispatch Radius, Forward-Looking, etc [Yan et al., 2018]. In addition, dynamic pricing algorithms such as Surge Pricing and Steady-State Models were also introduced. With regard to dynamic pricing, a key issue is the market failure, referred to as the "Wild Goose Chase" (WGC), which was briefly anticipated by Arnott in earlier discussions of radio taxi systems and has since re-emerged in ride-hailing services. Castillo et al. (2024) revealed that ride-hailing platforms such as Uber and Lyft are susceptible to a negative feedback loop of these matching failures during periods of high demand [Castillo et al., 2024]. Their study argued that surge pricing can effectively suppress demand and mitigate matching failures. They also suggested that surge pricing can slightly increase prices during peak hours while significantly reducing them during off-peak hours. Cashore et al. (2022) held a similar point and concluded that dynamic pricing provides robust equilibria in ridesharing networks [Cashore et al., 2022]. These studies motivate the consideration of the purpose of dynamic pricing, that is, the efforts made by the platform to balance the ever-changing, unpredictable, and difficult-to-balance supply and demand dynamics in the market.

Since the pricing mechanism used by ride-hailing platforms is designed to respond to fluctuations in demand and supply, a natural question arises: what factors influence needs and pools in this context? While several factors may seem intuitive based on everyday experience, empirical studies provide solid validation to support these beliefs. The first key element is substitution across transport methods. Natasha Vajravelu (2019) utilized the 2017 SAS Open as a natural experiment to specify public transit as a substitute for both ride-hailing and taxi services, with a remarkably more emphatic substitution effect on ride-hailing use [Vajravelu, 2019]. The study also discovered that this substitution is more noticeable on weekdays than weekends, implying there exists heterogeneity based on time and a more powerful effect on commute. While there is wide agreement that public transit alternatives exist for ride-hailing, the relationship between ride-hailing and traditional taxi services is more subtle. On the one hand, ride-hailing is often seen as a threat to the taxi industry, as evidenced by widespread reports of falling taxi ridership. Judd Cramer et al. (2016) argued that this disturbance derives from greater efficiency: higher capacity utilization allows ride-hailing to exceed taxis in operational terms [Cramer and Krueger, 2016]. Yet, the competition between these services is just one side of the story. Lisa Rayle et al. (2016) uncovered that while taxis and Uber share some overlap in user choices, ride-hailing also plays a role in supplementing travel vacancies and creating new demand [Rayle et al., 2016]. The study spotlighted two distinct behavioral mechanisms through which ride-hailing alters transportation decisions: cream-skimming, where existing modes are replaced, and gap-filling, where new travel demand is created that would not have occurred otherwise. Beyond substitution and behavioral change, the weather also plays a role in shaping the market. Abel Brodeur and Kerry Nield (2018) used rain in New York as an exogenous demand shock to compare ride-hailing and taxi responses under matching weather conditions [Brodeur and Nield, 2018]. They discovered that surge pricing incentivizes greater ride-hailing supply during adverse weather, while taxi drivers tend to reduce their service provision. Finally, temporal heterogeneity also matters, specifically across days of the week. James O. Huff and Susan Hanson (1986) observed that while travel behavior stays regular in the short term, there are notable structural variations across weekdays [Huff and Hanson, 1986]. Based on this insight, Axhausen et al. (2002) exhibited that weekday travel patterns are more regular and predictable, making them serve for modeling goals well [Axhausen et al., 2002]. Therefore, weekdays should be treated as fixed effects in empirical models rather than as random variations.

Looking back at earlier studies, my research mainly fills the hole in the use of natural experiments in empirical studies of platform algorithm price discrimination. Although there have been many studies on algorithm discrimination on ride-hailing platforms in the past, most of them used field studies (e.g. [Chang et al., 2022], [Ge et al., 2016]). The few studies that use natural experiments lack solid causal identification (e.g. [Pandey and Caliskan, 2021]). Exogenous natural experiments and massive real-world data have greatly improved the reliability of my research. In addition, this study also expands the application practice of the Diff-in-RDiT identification framework, which can provide a reference for hypothesis verification and causal identification for subsequent research.

3 Policy Background

Some critics of ride-hailing concentrated on its negative influences on transport systems. The primary debates are that ride-hailing is guilty of increased congestion: (1) adding vacant vehicle miles [Cramer and Krueger, 2016, Henao and Marshall, 2019]; (2) replacing for more efficient modes of transportation [Clewlow and Mishra, 2017, Rayle et al., 2016]; and (3) causing trips that otherwise would not have been created [Clewlow and Mishra, 2017, Rayle et al., 2016]. Further concerns have also arisen. John Manuel Barrios et al. claimed that the introduction of ride-hailing was associated with an approximate 3% growth in the number of fatalities and fatal accidents, consistent with the picture that ride-hailing increases congestion and overall roadway usage [Barrios et al., 2019].

A report[8] from the Chicago Government noted that from 2015 to 2018, ride-hailing trips increased from 27.6 million to 102.5 million. During the same period, total passenger mileage rose from 136 million miles to 603 million miles, representing a 344% increase. The report also highlighted that total mileage, including "deadheading" (i.e., empty driving and waiting for passengers), was even higher and ascribed no direct transport value. Ride-hailing trips have also strained traffic in already congested areas, mainly in high-traffic zones inc In the city center, the density of ride-hailing cars per foot of roadway is 21.7 times the citywide average. During peak hours, ride-hailing trips occupy approximately 26 miles of road in the downtown area, equal to six full lanes of Michigan Avenue.

In response, the City of Chicago implemented a policy that used taxation as a tool to alleviate traffic congestion across the city. On January 6, 2020, a congestion tax targeting ride-hailing services such as Uber, Lyft, and Via was implemented in the City of Chicago[9]. This policy aimed to mitigate increasing traffic congestion, particularly in the downtown area, and can be seen as a response to growing concerns that ride-hailing services exacerbate urban congestion. Because this tax policy applied uniformly across all groups and came into effect at a clearly defined date, it provides a hard-won quasi-experimental opportunity for causal inference.

This congestion tax policy introduced a tiered surcharge scheme based on both the time of day and the start and end locations of trips, with higher rates imposed for single-passenger rides during peak hours and for trips involving downtown Chicago (the Loop). In detail, trips that start or end in the Loop during peak hours (6 AM to 10 PM on weekdays) charge an additional surcharge of $1.75 per trip for single rides and $0.60 per

[8] Transportation Network Provider and Congestion in the City of Chicago: https://www.chicago.gov/content/dam/city/depts/mayor/Press%20Room/Press%20Releases/2019/October/TNPCongestionReport.pdf

[9] The Report of the Congestion Tax in the City of Chicago: https://www.chicago.gov/city/en/depts/mayor/press_room/press_releases/2019/october/NewRegulationsEaseTraffic.html

trip for shared rides. For trips outside these hours, the tax on all Transportation Network Provider (TNP) trips increased from \$0.72 to \$1.25 for single rides and from \$0.60 to \$0.65 for shared rides. In addition, the government implemented a \$5.00 surcharge on ride-hailing trips that start or end at designated locations, including O'Hare International Airport, Midway International Airport, Navy Pier, and McCormick Place. To confirm that the policy altered the price structure of ride-hailing services, I analyzed trips of each type in both 2019 (the placebo year) and 2020 (the treatment year). The components of ride-hailing prices exhibit changes that align with the policy implementation (see Appendix from Figure A1 to Figure A6).

4 Theoretical Framework: Price Discrimination in the Algorithmic Era

4.1 Classical Models of Price Discrimination

The classic book The Theory of Industrial Organization by Jean Tirole defines three degrees of price discrimination typically observed in the real world [Tirole, 1988]. First-degree price discrimination, where firms set prices specific to individuals and extract all consumer surplus, is considered unrealistic and hard to execute in markets. Second-degree price discrimination, by contrast, involves offering a menu of pricing options, allowing firms to yield self-selection among customers. This approach is more feasible and is generally applied in practice. For example, via the menu of ride-hailing service tiers offered by Uber and Lyft, these companies capture the paying willingness of consumers. Nevertheless, second-degree price discrimination is not the focus of this paper and this paper will not further discuss it. Third-degree price discrimination happens when firms can segment consumers into groups based on detectable features and charge each group a distinct price according to their demand elasticity. Third-degree price discrimination can also be achievable in the real world and is my primary perspective in the subsequent discussion.

Several adjustments must be considered when investigating price discrimination in the context of contemporary digital platforms, including ride-hailing services. With the improvement of big data technologies, customers' data is now routinely collected and investigated by companies. The idea of "letting data tell the story" is central to data science, and it has fundamentally altered the understanding of how price discrimination is executed in practice. The first important shift is that first-degree price discrimination can now be roughly enforced. Firms can use consumer purchasing history and personal characteristics to construct user profiles. Based on these profiles, firms can evaluate each consumer's willingness to pay and set personalized prices accordingly.

Both the legal system and the public have expressed concerns about the potential abuse of personal information. As a result, some regions have implemented restrictions on its use. In 2018, the European Union enacted the General Data Protection Regulation (GDPR), which requires pseudonymization as a standard process for storing personal data, offering protection against misuse to some extent. After that, in 2022, the American Privacy Rights Act (APRA) was proposed in the United States, similarly aiming to safeguard personal data from abuse. These developments have prompted renewed reflection, particularly on first-degree price discrimination. In short, first-degree price discrimination through direct access to personal information faces increasing challenges and has been restricted to some extent.

While first-degree price discrimination faces increasing challenges, third-degree price discrimination remains feasible in many real-world contexts. Suppose a platform faces two consumer groups, A and B, each with distinct demand functions, $Q_A(p)$ and $Q_B(p)$, respectively. Assuming a constant marginal cost and no loss of generality, the platform sets group-specific prices to maximize its total profit as follows:

$$\arg\max_{p_A, p_B} \pi = (p_A - c)Q_A(p_A) + (p_B - c)Q_B(p_B)$$

As a result, the firm can apply a markup given by $\frac{p_i - c}{p_i} = \frac{1}{\varepsilon_i}, \quad i \in \{A, B\}$ for each group, where ε_i represents the price elasticity of demand for group i.

This instance supposes that the platform quite observes each customer's group membership, a strong presumption that rarely holds in markets. In reality, specifically in the ride-hailing industry, the use of personal data has been prohibited by law. Even in the absence of legal constraints, platforms commonly do not observe individual-level economic characteristics such as earnings or price sensitivity. Yet, they often observe correlated signals, such as a user's location or community, which may serve as proxies for unobservable consumer features. In such scenes, the framework of Bayesian screening provides a more realistic and flexible approach.

4.2 Proxy-Based Bayesian Pricing and Algorithmic Screening

In many situations, platforms can only observe proxy characteristics rather than directly identify consumers' types. In such cases, the Bayesian screening framework provides a better description. In the context of this paper, ride-hailing platforms cannot directly observe each consumer's type, but they can rely on correlated signals—such as neighborhood-level economic indicators—to infer price sensitivity. Let θ denote the consumer type, and let Z denote observable community-level socioeconomic characteristics. The platform can then assume a conditional distribution $\theta \sim F(\theta|Z)$, and solve the following profit maximization problem:

$$\arg\max_{p(Z)} \pi(Z) = p(Z)(1 - F(p(Z)|Z))$$

That is, platforms determine prices based on associated socioeconomic characteristics. The probability that a consumer is willing to pay the price p, given Z, is $1 - F(p|Z)$. This creates the potential for algorithmic, proxy-based price discrimination, where pricing disparities emerge not from individual preferences, but from community-level signals.

The key difference between the proxy-based Bayesian pricing mechanism and the classic third-degree price discrimination lies in the information structure. In third-degree discrimination, the firm sets group-specific prices because it knows who belongs to which groups. In Bayesian screening, the firm only sees signals Z, and must set $p(Z)$ based on inferred type distributions.

The theoretical framework outlined above directly motivates my empirical approach. If a platform employs a proxy-based pricing strategy, I expect heterogeneous pricing adjustments in response to a uniform policy shock, such as the 2020 congestion tax in Chicago.

5 Data and Variable Construction

Since the research goal determines the data needs, this study first clarifies the objective before describing the datasets. The goal of this study is to identify whether ride-hailing pricing reflects socioeconomic-based discrimination. Ideally, if access to trip records of a group of well-known individuals were possible, that is, if trip records could be collected, I could obtain the socioeconomic backgrounds of each individual, and I also can match each trip record with the consumer, I can easily answer the two questions. Due to privacy regulations, individual-level data are unavailable. Additionally, I also need to consider whether there exists any confounders affecting the observability of the interested dependent. Hence, I control for observable confounders to reduce omitted variable bias. If I can collect the main confounders, I can leave less randomness in the model.

In the next subsections, I will discuss the trip records I can obtain, the socioeconomic background information I can collect, and the observable control variables. I also detail the data mapping process used to construct the indicators for empirical analysis.

5.1 Trip-Level Ride-Hailing Data

The Chicago Data Portal provides ride-hailing traveler data[10] from November 2018, containing a wealth of records used to analyze the development of ride-hailing services and offering insights into how ride-hailing plat-

[10]Transportation Network Providers Trips (2018-2022): https://data.cityofchicago.org/Transportation/Transportation-Network-Providers-Trips-2018-2022-/m6dm-c72p/about_data

forms determine pricing. These travel records are reported by ride-hailing companies operating in Chicago and are updated quarterly. Some recorded features were added from 2023[11], but most have remained the same since 2018. Since these new features do not significantly impact this research, they will not be further discussed in this paper.

The dataset is organized at the travel level, where each record corresponds to a unique trip. Per record contains fare, tips, additional charges, and the total trip cost, furnishing a comprehensive set of price signals. Moreover, the dataset captures trip time and trip distance, key determinants in pricing. Additionally, the dataset documents both the trip start time and end time, permitting the classification of distinct travel time spans. It also provides pickup and drop-off location information for matching the trip records with the socioeconomic features. The dataset also includes attributes indicating whether a trip is shared and whether it was successfully matched with other passengers. The shared trip indicator implies that the ride may include multiple stops and destinations, a characteristic feature of the sharing economy.

Some limitations of the dataset must be recognized. The most influential weakness is the anonymization of trips[12]. Although ride-hailing enterprises do not provide passengers' data to the administration, certain techniques could potentially use precise trip details, such as exact locations and trip times combined with other data sources to reidentify passengers. In compliance with personal privacy protection regulations, the Chicago government uses data processing procedures to anonymize personal information. To safeguard passenger privacy, the dataset rounds all trip times to the nearest 15-minute interval, and latitude and longitude coordinates are not provided. Rather, trips are reported using the census tract center of the pickup and drop-off locations. Additionally, fares are rounded to the nearest \$2.50, and tips are rounded to the nearest \$1.00. These data preprocessing steps introduce challenges for research and the following subsection discusses how these constraints are addressed.

This research selected a subset of attributes from the dataset while excluding others. Fare and additional charges are chosen as price indicators, whereas tips and the total trip charge are excluded. Fare and additional charges reflect the fundamental pricing mechanism, dynamic pricing adjustments, and policy effects such as tax changes, while tips are individually determined, reflecting consumer behavior differences, which cannot be utilized since the dataset does not provide individual-level data. Trip time and trip distance are incorporated as controllers in the model, as they mainly decide trip prices. The start and end times of trips are also considered, as they equip knowledge to distinguish between weekdays and non-weekdays and to categorize rush hours versus non-rush hours. I distinguish between weekdays and non-weekdays, as well as rush hours and non-rush hours, because they exhibit different demand patterns (see Appendix Figure A7) and are associated with different pricing structures. This study incorporated pickup and drop-off location features to connect with census datasets, enabling the analysis of the socio-economic background of the consumers of each trip. Additionally, all shared trip records are excluded, as multiple stops and destinations introduce unnecessary confounding factors and may follow a different pricing mechanism.

5.2 Socioeconomic Data and Spatial Mapping

Since this study examines whether ride-hailing platforms apply price discrimination based on passengers' socioeconomic backgrounds, socioeconomic data are another key dataset that must be collected. As previously stated, the Chicago Data Portal does not provide passengers' personal information; instead, it aggregates data using attributes such as the pickup census tract and pickup community code. Consequently, this study cannot analyze trips at the individual level but rather at the group level.

At the aggregated level, two different spatial units can be chosen for analysis: the census tract level and the community level. Both of the units are coded within the trip-level dataset, and it is known that there are 77 communities in the City of Chicago, while there are 800 census tracts in the same area. One approach is to study at the census tract level, while another option is to analyze data at the community level. These units provide different levels of time granularity for analysis. The United States Census Bureau provides

[11]Transportation Network Providers Trips (2023-2024): https://data.cityofchicago.org/Transportation/Transportation-Network-Providers-Trips-2023-2024-/n26f-ihde/about_data
[12]Personal Privacy Guidelines: https://data.cityofchicago.org/stories/s/82d7-i4i2

the American Community Survey (ACS)[13], which includes socioeconomic data at the census tract level, covering the city of Chicago. The ACS is an ongoing survey that provides annual data on social, economic, housing, and demographic characteristics of the U.S. population. Apart from the annual survey, the ACS offers non-overlapping 5-year estimates, which deliver a more stable population estimate. In addition, the Chicago Metropolitan Agency for Planning (CMAP)[14] also provides supplementary data covering all 77 Chicago community areas, including public demographic, housing, and employment.

Although both datasets provide attributes relevant to this study, in this research, I use the 5-year ACS dataset, which supports a census tract-level analysis. Upon inspecting the dataset, I focus on income-related indicators at the community level, as they reflect the socioeconomic background and overall economic development of each area (see Appendix Figure A8 and Figure A9). Specifically, the variables include median household income, per capita income, income bracket counts, poverty rate, participation in food stamp (SNAP) programs, and the share of households receiving social security income. The data itself does not directly reveal the heterogeneity in community classifications. Although economic indicators clearly differ across communities, it is not straightforward to define a classification criterion based solely on these variables. Therefore, additional data manipulation is required to derive the variables used in this study.

It must be acknowledged that the study I conducted was unable to track individual users due to privacy-preserving data anonymization mechanisms that prohibit access to personally identifiable information, and was necessarily aggregated at the census tract level, using community-level socioeconomic indicators as proxies for individual characteristics. However, this limitation is not new; this approach is widely used in public health, epidemiology, and social sciences and has been repeatedly validated in methodology (e.g. [Currie and Yelowitz, 2000], [Bayer et al., 2007], [Chetty et al., 2016]). This is always an accepted compromise when direct individual-level socioeconomic data are unavailable. Although such treatments have limitations such as the ecological fallacy that have been recognized in past studies, this area-based measurement approach can still provide valuable approximations, especially when the analysis focuses on structural or group-level patterns rather than individual-level inferences.

5.3 Control Variables: Transportation Substitutes and Weather

Ride-hailing companies such as Uber and Lyft claim that pricing is decided by supply and demand dynamics. Thus, factors affecting these dynamics should be assessed. Some controllers can be generated from the TNP dataset, like rush hours, as they significantly affect how demand fluctuates throughout the day.

To control for confounding factors, this study considers substitute goods and weather conditions. The Chicago Data Portal provides traditional taxi data[15] from 2013, which follows a similar feature layout as the Transportation Network Providers (TNP) Trips dataset described in the previous section. The traditional taxi trips dataset is particularly important, as a body of literature have identified traditional taxis and ride-hailing services as substitute goods with overlapping consumers [Cramer and Krueger, 2016]. Additionally, this study collects "L" rails flow data[16] and bus flow data[17] from the Chicago Data Portal, which includes bus and rail services provided by the CTA, with records dating back to 2001. Chicago's "L" rail system is the city's urban rapid transit network[18]. Operated by the CTA, it serves downtown Chicago and extends into the surrounding suburbs. The city's bus system, also operated by the CTA, is the second-largest public transportation network in the United States, covering the city and more than 40 surrounding suburbs. The system includes 127 bus routes, providing over 18,500 trips per day and serving more than 10,500 bus stops[19]. The "L" system and the CTA bus network are the major modes of public transportation in the City

[13]American Community Survey 5-Year Data (2009-2023): https://www.census.gov/data/developers/data-sets/acs-5year.html

[14]The Chicago Metropolitan Agency for Planning's (CMAP) Community Data: https://cmap.illinois.gov/data/community-data-snapshots/

[15]Taxi Trips (2013-2023): https://data.cityofchicago.org/Transportation/Taxi-Trips-2013-2023-/wrvz-psew/about_data

[16]"L" Station Entries Ridership: https://data.cityofchicago.org/Transportation/CTA-Ridership-L-Station-Entries-Daily-Totals/5neh-572f/about_data

[17]Bus Routes Ridership: https://data.cityofchicago.org/Transportation/CTA-Ridership-Bus-Routes-Daily-Totals-by-Route/jyb9-n7fm/about_data

[18]"L" System History: https://www.chicago-l.org/history

[19]CTA Bus System Service Overview: https://www.transitchicago.com/facts

of Chicago and hence serve as proxies for public transit access in this study.

Obtaining weather data is easy to some extent; the National Centers for Environmental Information (NCEI)[20] provides comprehensive weather records for the city of Chicago, including average wind speed, precipitation, snowfall, and average temperature.

Unlike the taxi data, public transportation and weather data are recorded at the daily level. This limitation constrains the time granularity, preventing the study from conducting a comprehensive supply-demand analysis within time intervals shorter than a day. As a result, an underlying assumption must be made: weather conditions affect the supply and demand for ride-hailing services at the daily level, as does public transportation. Fortunately, the taxi data maintains the same level of granularity as the TNP data, allowing for more precise analysis at the hourly level.

5.4 Data Integration and Preprocessing Strategy

After collecting and analyzing the data composition as described in the previous section, several additional preprocessing steps are required. An essential but straightforward step is to clean the data and restrict its scope. Specifically, I dropped all records in which trips start or end outside the city of Chicago, ensuring that the analysis focuses exclusively on trips within city boundaries. I retained only the trips that started between 6 a.m. and 10 p.m. on workdays, as the congestion tax specifically targeted these time frames. Given the time and distance of trips, I dropped trips with time or distance lower than the 5% lower bounder and higher than the 95% upper bounder to remove the influence of extreme trips.

I calculated the sum of fare and additional charge, denoted the result as price. This is because the fare was mainly determined by the time and distance of the trip, and the additional charge including the taxes, fees, and any other charges for the trip. A combination of fare and additional charge can reflect the change in tax as well as the pricing algorithm.

To capture calendar-related cyclical effects, I generated the day of the week from each trip's date and transformed it into a set of dummy variables. Particularly, I constructed four dummy variables representing Tuesday, Wednesday, Thursday, and Friday, based on the observed pattern of daily cyclical instabilities.

Given that the congestion tax includes a special zone fee and imposes a downtown surcharge compared to other areas, I classified trips into three groups based on their pickup and drop-off locations at the community level. Trips to, from, or within the special zone are coded as 2; trips with either pickup or drop-off in the downtown area are coded as 1; and all other trips are coded as 0. Similar to the day-of-week process, I convert these classifications into a set of dummy variables for regression analysis.

Although this study collected daily ridership data for both the "L" system and the CTA bus system, these data are not directly linked to census tracts. To address this limitation, I additionally collected station location data[21] for the "L" system and spatial route data[22] for the CTA bus system from the Chicago Data Portal. I used geographic information system techniques to integrate these spatial datasets with the ridership data. For the "L" system, I mapped station entries to their geographic coordinates. I then merged the processed daily ridership data with the American Community Survey (ACS) dataset to estimate daily "L" system ridership at the census tract level. For the bus system, the process was more complex. Since the Chicago Data Portal only provides ridership data at the route level, I adopted an indirect estimation approach. Specifically, I first calculated the proportion of each bus route that passes through census tracts and multiplied these proportions to distribute the daily route-level ridership. Finally, I merged the processed bus ridership data with the ACS dataset to estimate daily bus system ridership by census tract. One additional step applied to the two ridership datasets is that I normalized the ridership data to remove unit and scale differences, in order to avoid the clear order-of-magnitude gap between raw ridership and price data.

[20] National Centers for Environmental Information: https://www.ncei.noaa.gov

[21] List of "L" Stops: https://data.cityofchicago.org/Transportation/CTA-System-Information-List-of-L-Stops/8pix-ypme/about_data

[22] CTA Bus Routes: https://data.cityofchicago.org/Transportation/CTA-Bus-Routes/6uva-a5ei/about_data

To ensure the internal validity of my quasi-experimental framework, I consider appropriate temporal granularities. Both the ride-hailing trip data and the taxi trip data are recorded to the nearest 15 minutes. The weather and public transportation data are recorded at the daily level, while census tract socioeconomic data are estimated at five-year intervals. I assume that weather conditions are uniform across all trips within the same hour of the day, and that public transportation competes with ride-hailing services at the daily level. Taxi trip data are aggregated to the 15-minute level, and I construct taxi flow for each 15-minute interval within each census tract to match the ride-hailing (TNP) trips. This serves as a direct substitute for ride-hailing services during each time period and within each pickup census tract. I also assume that socioeconomic background remains stable throughout the quasi-experimental period.

Lastly, to ensure that multicollinearity does not bias my research results, I compute Variance Inflation Factors (VIFs) for all potential control variables. In brief, the VIF quantifies how much the variance of a coefficient increases due to linear dependence among regressors. Given the results, it shows that almost all of the control variables have VIFs below 5, indicating low collinearity (see Appendix Table B.1 and Table B.2). Although the VIFs for both "L" rail system and bus are relatively high (above 10), suggesting a strong correlation between rail and bus usage, I retained both variables in the model. This decision is based on their distinct conceptual roles: while both capture aspects of public transportation demand, they reflect different transportation modes and usage patterns. Including both allows this study to more comprehensively control for variations in overall transit behavior, minimizing omitted variable bias, even at the cost of tolerating some degree of multicollinearity.

6 Empirical Strategy: A Diff-in-RDiT Approach

A simplistic approach might suggest testing the correlation between prices and income levels; however, such tests are insufficient to identify causal effects. The determination of fares is complex and may contain many different confounding factors. This makes it impossible to clearly determine whether the correlation obtained by statistical regression comes from the price discrimination of the platform. Although such a regression can use control variables to absorb some fluctuations caused by errors, even with a relatively clear correlation, it is still impossible to clearly determine the direction of causality. Therefore, the study needed to design a causal inference framework to isolate the platform's price discrimination from numerous confounding factors. In causal inference, researchers often want to find an effective exogenous shock to isolate confounding variables. Fortunately, an exogenous shock was observed, namely the implementation of the ride-hailing congestion tax policy in the city of Chicago. While the tax does not explicitly target low-income or minority communities, the black-box nature of ride-hailing algorithms may result in differential pricing adjustments depending on community characteristics, which constitutes the focus of this study. This natural experiment does have a volume of data that can be analyzed, as I have introduced, and it can be used to design a robust identification framework for causal inference.

6.1 Identification Framework

6.1.1 Baseline Framework: Regression Discontinuity in Time (RDiT)

To identify the effect of the tax policy on pricing outcomes and whether these effects vary by community-level socioeconomic status, I first use the sharp timing of the policy implementation using a Regression Discontinuity in Time (RDiT) framework. As a type of regression discontinuity design (RDD), my study draws on several insights from earlier research. For instance, Hansen (2015) emphasized that a flexible RDD relies on the assumption that the identification points are differentiated by physical indicators that cannot be manipulated by humans [Hansen, 2015]. As the saying goes, "time and tide wait for no man", and it is credible to believe that time cannot be manipulated, nor can individuals systematically alter their behavior precisely at the cutoff. This approach compares observations before and after the policy threshold under the identifying assumption that, in the absence of treatment, trip prices would have evolved smoothly over time.

To construct an efficient estimation model, it is necessary to confirm which controllers should be included in the model. Based on the previous discussion, I included trip-level factors (e.g., time, distance, rush hour), temporal indicators (e.g., days since policy implementation, day of the week), external conditions (e.g., weather), availability of substitutes (e.g., public transportation and taxis), and policy-affected area types, as each of these factors may directly or indirectly influence pricing. Specifically, time, distance, and rush hours are key determinants of ride-hailing prices and have been discussed in the literature. Weather affects both supply and demand by affecting travel decisions and driver service willingness, particularly during extreme conditions that may reduce travel. Substitutes, including public transportation (e.g., bus and "L" system) and taxis, affect pricing through competition, as travelers may switch modes relying on their convenience and trip needs. These controllers help isolate the treatment effect by accounting for variation unrelated to the policy intervention.

6.1.2 Extended Framework: Difference-in-RDiT (Diff-in-RDiT)

While the RDiT design and parametric controllers above strengthen internal validity, they may not fully eliminate bias arising from cyclical patterns—such as weekday-weekend transitions or seasonal travel behaviors. That is although I have included the temporal indicators, there are temporal confounders with nonlinear or unobserved structures may bias the estimated treatment effect. To address this limitation, I incorporate ride-hailing trip data from December 8, 2018, to February 8, 2019, as a comparable control year. This timeframe matches the duration of the treatment window in my main analysis (from December 8, 2019, to February 8, 2019). The Pearson correlation coefficient between daily trip volumes across the two years needed to be computed to check the comparability of these two time windows. I chose January 7, 2019, as a hypothetical cutoff of the control year. The reason I did not choose the exact same date in 2019 is that I needed to account for the influence of weekdays and weekends. Therefore, this study selected January 7, 2019, a Monday, to match the weekday of the policy implementation date in 2020, which also fell on a Monday. At this stage, I need to estimate the difference before and after the hypothetical cutoff in 2019, following a process similar to the one used for the treatment year. Then, to account for cyclical effects, I can remove the estimated difference at the corresponding hypothetical cutoff in the control year from the estimated jump at the true cutoff in the treatment period. This difference-in-discontinuities method helps to isolate the causal impact of the policy from time-based confounding factors. Given the above analysis, this study constructs the following Diff-in-RDiT model:

$$
\begin{aligned}
Y_{it} =& \beta_0 + \beta_1 \cdot \text{Post}_t + \beta_2 \cdot \text{TreatYear}_i + \beta_3 \cdot (\text{Post}_t \cdot \text{TreatYear}_i) \\
& + f(\text{TFC}_t) + f(\text{TFC}_t) \cdot \text{Post}_t + f(\text{TFC}_t) \cdot \text{TreatYear}_i + f(\text{TFC}_t) \cdot (\text{Post}_t \cdot \text{TreatYear}_i) \\
& + \gamma^\top X_{it} + u_{it}
\end{aligned}
$$

In this model, Y_{it} denotes the price of trip i at time t. Post_t is a binary indicator equal to 1 if the trip occurred after the policy implementation date, either the actual date in the treatment year (January 6, 2020) or the hypothetical date in the control year (January 7, 2019), and 0 otherwise. TFC_t stands for "time from cutoff", and centers the time variable at the cutoff date to capture temporal trends in pricing. For example, both January 4, 2020, and January 8, 2020, are coded as being 2 days from the actual policy cutoff. In addition to centralizing the time variable, the degree of polynomials used in the main model should also be clarified. I adopt a first-order linear function for $f(TFC_t)$ to capture trends around the cutoff. This method balances parsimony and interpretability while retaining the flexibility to extend to higher-order polynomials in future studies. X_{it} is a vector of controls that accounts for residual variation driven by other factors. These controls include trip duration, distance, rush hour indicators, weather conditions, transportation substitutes, day of the week, and policy-affected area types as previously discussed.

$$
Y_{it} = \beta_0 + \beta_1 \cdot \text{Post}_t + \beta_2 \cdot \text{TreatYear}_i + \beta_3 \cdot (\text{Post}_t \cdot \text{TreatYear}_i) + u_{it}
$$

To clarify the model, I begin with the simplified specification shown above. In this model, β_0 represents the baseline outcome for trips that occurred between December 8, 2018 and January 6, 2019; $\beta_0 + \beta_1$ represents outcomes for trips between January 7, 2019 and February 8, 2019; $\beta_0 + \beta_2$ corresponds to trips between

December 8, 2019 and January 5, 2020; and $\beta_0 + \beta_1 + \beta_2 + \beta_3$ represents outcomes for trips between January 6, 2020 and February 8, 2020.

The Difference-in-Differences (DiD) framework worked as the base of the final major model. Based on it, I introduce the running variable and integrate it into the original model. Interaction terms are used to capture fluctuations across the four distinct periods. Subsequently, I incorporate control variables and heterogeneity terms to construct the final model.

In this specification, β_3 is the coefficient of interest as it captures whether a post-policy price jump occurred in 2020, serving as evidence of the policy's effect. Since I am interested in whether there are differences among heterogeneous groups, I estimate the model across stratified subsamples to obtain a set of coefficient estimates. To classify census tracts into different groups, I applied a K-means clustering algorithm using the KNN-based aggregation method. The model was fitted using a set of socioeconomic features, including median household income, per capita income, income distribution, poverty rate, the percentage of households receiving food stamps/SNAP, and the percentage of households with social security income. Based on the clustering results, the census tracts were divided into three income groups, denoted as low-income mid-income and high-income census tracts (see Appendix Figure A10).

6.2 Key Identification Assumptions and Validity Tests

There is a consensus to make sure some assumptions are valid before using the model. Since the Diff-in-RDiT approach combines elements of both Difference-in-Differences and Regression Discontinuity, it is essential that the identifying assumptions of both methods are satisfied.

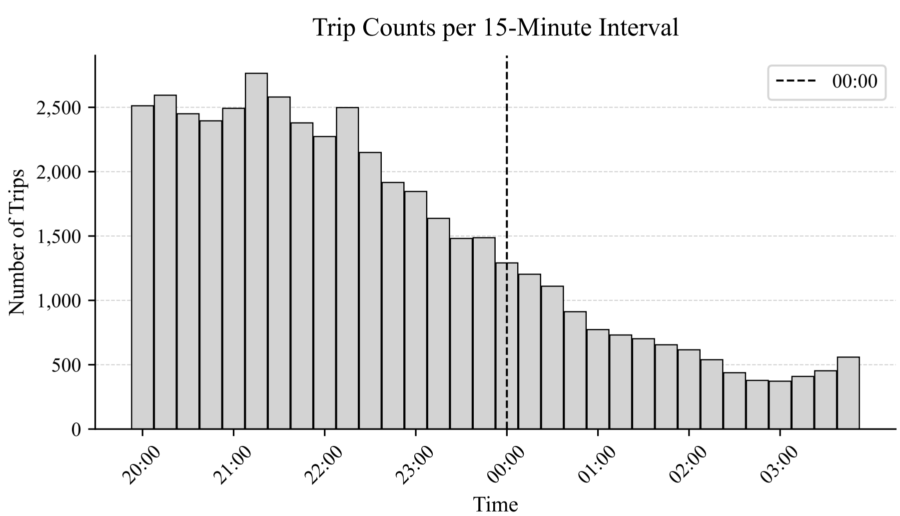

Figure 1: **Trip Counts per 15-Minute Interval Around Policy Cutoff.** This figure displays the distribution of ride-hailing trips in 15-minute intervals from 8:00 PM on January 5 to 4:00 AM on January 6, 2020. The black vertical dashed line at 00:00 marks the policy cutoff time when the congestion tax takes effect. The absence of a discontinuous jump at the cutoff suggests no evidence of strategic timing behavior by users, supporting the validity of the running variable in the Diff-in-RDiT design.

The RDiT framework relies on the assumption that, in the absence of treatment, the outcome variable, price, would evolve smoothly and continuously across the cutoff. An implicit assumption of the design is that the timing of treatment is exogenous to the behavior of riders and drivers. Given the nature of the policy announced by the city and implemented uniformly across the platform, I assume that individuals cannot precisely manipulate the timing of their trips to avoid or exploit the tax. Furthermore, I conducted the McCrary test, following a method similar to that used by Deshpande [Deshpande, 2016], on the raw trip data. The results, presented in Figure 1, indicate that there is no bunching around the cutoff point. As the distribution drops smoothly across the cutoff, I can reasonably conclude that there is no evidence of manipulation bias.

Another implication, any observed discontinuity in the outcome can be attributed to the causal effect of the policy intervention. While the full model specification is used for causal inference and subgroup analysis, I employ a simplified-form version in assumption checking. Specifically, the simplified model shown on below facilitates pre-treatment covariate balance checks, without conflating identification assumptions.

$$Y_{ij} = \beta_0 + \beta_1 \cdot \text{Post}_t + f(\text{TFC}_t) + f(\text{TFC}_t) \cdot \text{Post}_t + \gamma^\top X_{it} + u_{it}$$

I conducted pre-treatment covariate balance tests for two types of variables. For weather-related factors, I ran regressions without covariates. For substitute-related factors, I included day-of-week fixed effects as controls, given that commuter travel patterns exhibit calendar-based cyclicality. As previously discussed, weather data and public transit volume are recorded at the daily level, whereas taxi data share the same temporal granularity as ride-hailing data. Accordingly, I test the smoothness of these variables at different levels. Specifically, weather features and public transit volume are evaluated at the daily level, while taxi data are examined at the level of trip start time interacted with pickup census tract. This allows us to capture detailed information on substitute availability for ride-hailing services at specific times and locations.

I choose a bandwidth of 15 days without loss of generality, and the test results are presented in Table 1 and Table 2. According to Table 1, all four weather indicators vary smoothly across the cutoff, as their estimated slope coefficients are statistically insignificant. As shown in Table 2, public transit volume also changes smoothly across the cutoff. Taxi usage, on the other hand, exhibits a significant jump at the cutoff, suggesting the possibility of confounding. However, this does not pose a concern, as taxi volume is included as a control variable in the main specification.

Table 1: Smoothness Test of Weather Variables

	Average Temperature (°C)	Precipitation (mm)	Snowfall (mm)	Wind Speed (m/s)
Treatment Year				
Treatment	1.121	2.166	-11.253	-0.357
	(3.959)	(5.381)	(19.029)	(1.389)
Mean at cutoff	0.905	1.775	6.100	4.565
Controls	No	No	No	No
Observations	20	20	20	20
Placebo Year				
Placebo Treat	1.122	-1.267	-3.459	0.475
	(4.133)	(6.393)	(7.215)	(1.699)
Mean at cutoff	-1.255	2.985	2.850	4.360
Controls	No	No	No	No
Observations	20	20	20	20

Notes: This table reports the smoothness tests for weather variables around the policy cutoff. “Treatment” refers to the coefficient for the treatment indicator in the treatment year, while “Placebo Treat” refers to the same coefficient in the placebo year. No additional control variables are included in these regressions. Robust standard errors are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: Smoothness Test of Substitute Transportation Modes

	L Rail	Bus	Taxi
Treatment Year			
Treatment	0.262	0.167	0.132***
	(0.789)	(0.756)	(0.018)
Mean at cutoff	-0.185	-0.118	0.136
Controls	Yes	Yes	Yes
Observations	20	20	60483
Placebo Year			
Placebo Treat	0.658	0.832	0.000
	(0.919)	(0.942)	(0.000)
Mean at cutoff	-0.278	-0.299	0.000
Controls	Yes	Yes	Yes
Observations	20	20	57667

Notes: This table reports the smoothness tests for potential substitute transportation modes around the policy cutoff. Each column corresponds to a different mode ("L" Rail System, Bus, Taxi). "Treatment" refers to the coefficient for the treatment indicator in the treatment year, while "Placebo Treat" refers to the same coefficient in the placebo year. Robust standard errors are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

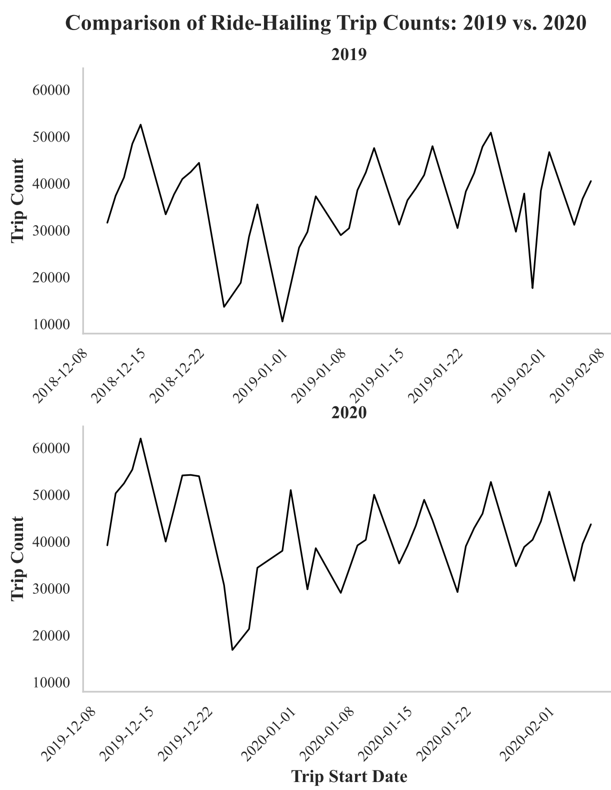

Figure 2: **TNP Trips Records Comparison.** These line charts compare the total number of ride-hailing trips recorded in the City of Chicago between two periods: December 8, 2018, to February 8, 2019 (top panel), and December 8, 2019, to February 8, 2020 (bottom panel). Each point represents the total daily trip counts. The charts illustrate seasonal patterns, fluctuations around the holidays, and potential differences across years.

To assess the parallel trends assumption, I plot the trip counts for the treatment and control years using a lag-adjusted calendar window, as shown in Figure 2. The key identifying assumption is that, absent the policy, discontinuities around the same calendar cutoff should be comparable across years. Although the travel volumes in the two time periods differ in magnitude, they exhibit similar peaks and troughs at corresponding points in time. I also conduct a quantitative correlation analysis. Based on correlation analysis, the optimal lag is −1. I find that the lag-adjusted Pearson correlation of price trends between treatment and control groups is 0.66, providing suggestive evidence of comparable pre-policy dynamics.

I also completed an event study following an approach similar to that of Miller [Miller et al., 2021]. In the event study, I use two weeks before the policy implementation (week –2) as the baseline period, construct dummy variables for each event week relative to the baseline, and interact them with the treatment year indicator to estimate the treatment effects over time. The model used in the event study is displayed below:

$$Y_{it} = \alpha + \lambda \cdot \text{Treat}_i + \sum_{w} \gamma_w \cdot \mathbf{1}\{\text{EventWeek}_i = w\} + \sum_{w \neq -2} \delta_w \left(\mathbf{1}\{\text{EventWeek}_i = w\} \times \text{Treat}_i\right) + \epsilon_{it}$$

The estimated coefficient for each event week can be interpreted as the difference in the treatment year relative to the control year, compared to the week before the event. The results show that while no significant differences are observed at week –3, a statistically significant positive difference emerges at week −1. This suggests the presence of some pre-treatment divergence, potentially indicating anticipatory effects or other confounding factors. Therefore, the strict parallel trends assumption may be partially violated. However, since the magnitude of pre-treatment differences is small compared to post-treatment effects, the deviation is minor and unlikely to materially affect the main findings. Forming the week of the event (week 0), prices grew significantly and remained elevated for the subsequent four weeks, showing that the policy had an immediate and continuous impact on the pricing of online ride-hailing services.

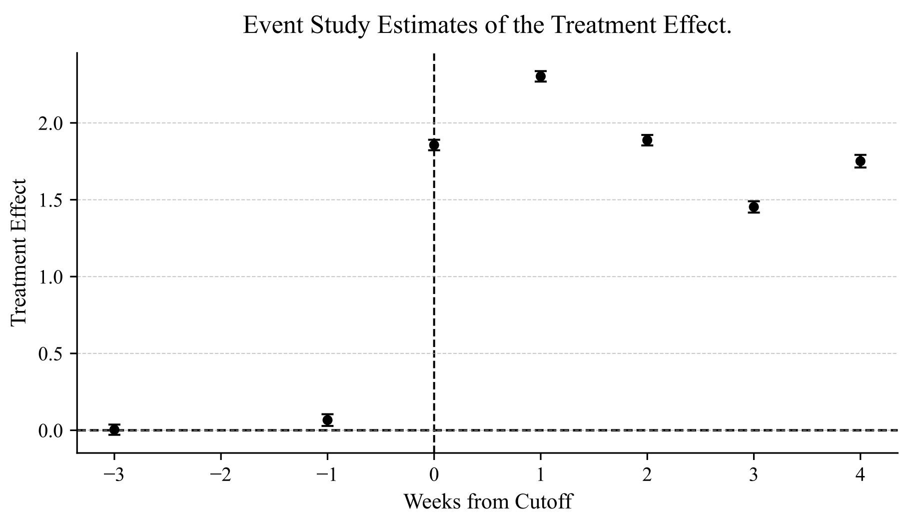

Figure 3: **Event Study Estimates of the Treatment Effect.** This figure shows the estimated dynamic treatment effects of the policy over time relative to the cutoff date. Each point represents the estimated difference-in-discontinuities effect for a given week, comparing 2020 (treated year) with 2019 (control year). The omitted week (-1) serves as the reference period. The vertical dashed line at week 0 marks the policy implementation date.

7 Empirical Findings and Heterogeneity Analysis

Now, fortunately, the Chicago ride-hailing congestion tax policy provides an exogenous natural experiment, and I have obtained a wealth of databases and provided a clear and reliable causal identification framework.

Therefore, the corresponding empirical results can be provided with relative confidence. Next, I will first present my results and rationale for community grouping, followed by the results of the aggregated group-level analysis. Then, I will display the causal identification results based on each travel record, and finally, I will demonstrate the robustness of the method.

7.1 Income-Based Community Clustering

Building on the spatial polarization literature that documents growing socioeconomic segregation in Chicago [Smith et al., 2021], I classify communities into income-based clusters to examine how algorithmic pricing may respond to this persistent urban stratification. I cluster census tracts into distinct socioeconomic groups using unsupervised learning. Before clustering, I reduce the dimensionality of the socioeconomic space using Principal Component Analysis (PCA), retaining the first two principal components that capture the majority of variance. As shown on Table A3, Principal Component 1 is strongly negatively loaded on median and per capita income, and positively on poverty-related indicators such as poverty rate and SNAP participation, indicating that it captures an "affluence-poverty" spectrum. The optimal number of clusters, $K = 3$, is selected using a combination of the Elbow Method and Silhouette Score criteria. After clustering, I revert to the cluster centroids as the basis for interpretation. Smith et al. [Smith et al., 2021] describe a growing divergence between affluent neighborhoods such as Lincoln Park and low-income areas in the South and West sides—patterns that align closely with my identified cluster centroids (see Appendix Figure A10).

Table 3: Cluster Centers Based on Census Tract Socioeconomic Indicators

	Low Income	Middle Income	High Income
Median Household Income	33951.60	62719.75	110398.95
Per Capita Income	19335.87	31742.69	59679.78
Income < 10k	89.55	88.58	26.87
Income 10k–15k	44.78	45.94	11.36
Income 15k–25k	106.39	137.38	35.49
Income 25k–35k	81.09	140.94	41.51
Income 35k–50k	92.66	200.04	70.76
Income 50k–75k	99.98	307.43	143.08
Income 75k–100k	51.01	224.53	133.75
Income 100k–150k	48.31	271.46	222.35
Income 150k–200k	8.61	97.28	140.53
Income > 200k	6.05	63.92	223.62
Poverty Rate (%)	20.73	7.81	2.27
Households on SNAP	247.05	195.11	24.32
Households with Social Security	218.71	467.00	226.04

Notes: This table reports the inverse-transformed cluster centers derived from K-means clustering on standardized socioeconomic variables.

Based on the characteristics of the three cluster centers reported in Table 3, I interpret the resulting clusters as representing low-income, middle-income, and high-income communities. Table 3 presents the socioeconomic characteristics of the three cluster centroids, which represent each of the clustered groups. As shown in the table, the low-income group exhibits a median household income of approximately \$33,951, with over 20% of residents living below the poverty line. In contrast, the high-income group features a median income of \$110,399, a poverty rate of just 2.27%, and substantially greater shares of households earning over \$150k. Moreover, the middle-income cluster falls in between, with a median income of around \$62,720 and a poverty rate of 7.81%. The allocation of these indicators is very consistent with the intuitive impression, which illustrates the effectiveness of the application of the unsupervised machine learning method.

7.2 Aggregated-Level Treatment Effects

After classifying the 800 census tracts in the City of Chicago into three income-level categories, I first perform regressions at the aggregated group level. Specifically, I aggregate the trips by calculating the average price within each group, which serves as the dependent variable. For the sake of simplicity, these preliminary regressions do not include control variables and assume a linear trend. The results are presented in Figure

4. First of all, it can be verified that the undertaking of the policy (congestion tax) brought a significant increase in prices, because in the treatment year (2020), at cutoff = 0, there is a clear jump in prices in all four figures. Further findings include that the price increase is not caused by natural factors such as seasonality and holidays, but is related to the implementation of real policies. This is because in the placebo year (2019), the price trend near the placebo cutoff is stable and there is no jump. In addition, before the cutoff, especially in the range of -30 to 0 days, the price trends of each group are relatively stable, and there is no sign of early change. The magnitude of the price jump (from before the policy to after the policy) of the low-, mid-, and high-income groups is slightly different, which initially suggested that there is significant heterogeneity in the policy effect among different income groups. To add some description, the R^2 of the pooled regression is approximately 0.94. For the income-specific groups, the R^2 values are approximately 0.92 for the low-income group, 0.93 for the middle-income group, and 0.94 for the high-income group.

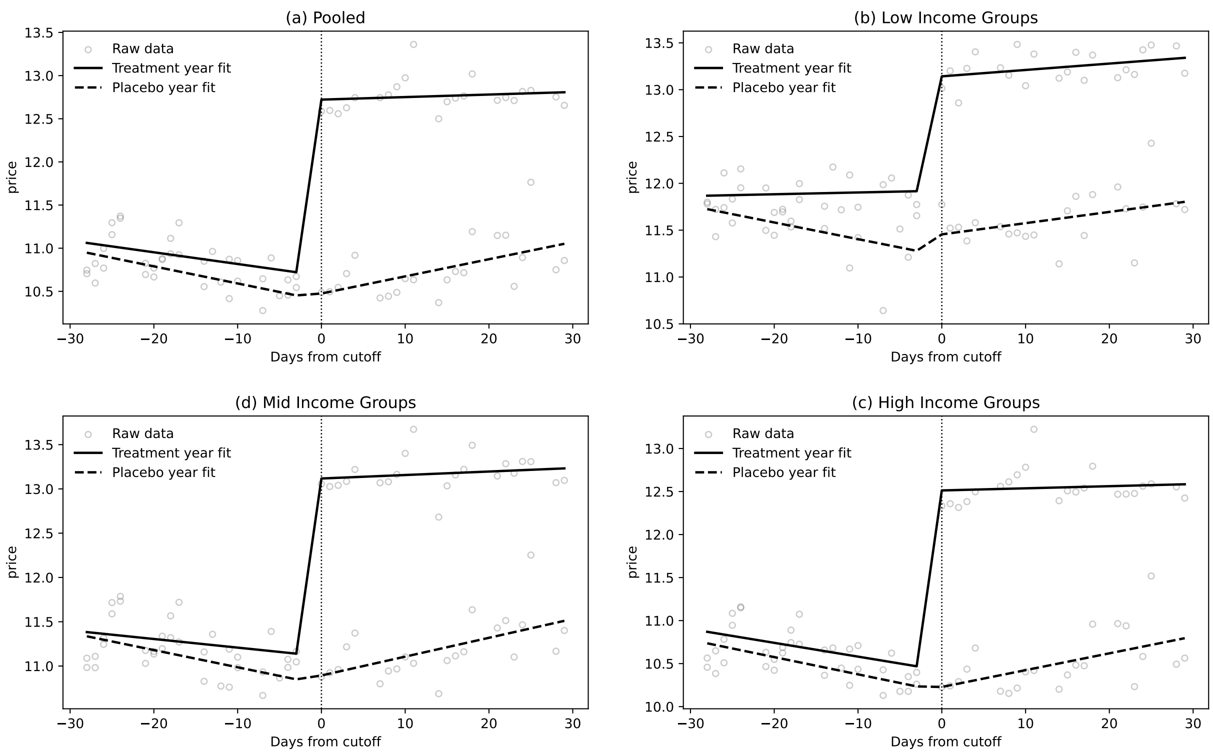

Figure 4: **Diff-in-RDiT Estimates Around the Congestion Tax Cutoff.** This figure presents the fitted values and raw data for the average trip price around the congestion tax policy cutoff, based on a Difference-in-Regression Discontinuity-in-Time (Dif-in-RDiT) design. Panel (a) displays the pooled results, while panels (b) to (d) report estimates separately for low-, middle-, and high-income groups. Solid lines indicate the treatment year fits, based on the actual policy implementation date (January 6, 2020), whereas dashed lines represent the placebo fits, assuming a hypothetical policy implementation date (January 7, 2019).

7.3 Individual-Level Treatment Effects

The aggregated group-level regressions are used solely for simplified visualization. To estimate treatment effects across groups, I ran the Diff-in-RDiT regression model using individual-level trip data, grouped by heterogeneity, with a 15-day bandwidth and a linear trend.

Table 4 presents the estimated coefficients for a set of control variables across different income groups. Table 4 offers the estimated coefficients for a set of control variables across different income groups. Overall, temperature and wind speed exhibit significant negative effects on the price of ride-hailing, showing that adverse weather conditions either reduce travel demand or increase travel costs. Yet, precipitation and snowfall display adverse effects, suggesting that people continue to commute under such conditions. Notably, the effects of temperature and precipitation deviate by income group: these estimates are statistically significant for middle- and high-income people, but insignificant for low-income travelers. Taxi usage is positively associated with the price of ride-hailing, especially among low-income groups, revealing that ride-hailing and taxis show more of a complementary effect. This finding supports the argument that ride-hailing functions as a supplementary mode of transportation [Rayle et al., 2016]. The connection between buses, the "L" system, and ride-hailing prices is highly heterogeneous across income communities. First, checked the bus system, and I found that in low-income communities, increased bus use is associated with significantly higher ride-hailing prices, potentially reflecting stronger demand rigidity due to limited transportation options or time-sensitive travel needs. As a comparison, middle-income areas display an obvious substitution effect, where increased bus use corresponds to lower ride-hailing prices. In high-income communities, the weak positive relationship may reflect overlapping demand surges during major events rather than a direct substitution effect. When I turn my attention to the "L" system, I can also find clear between-group heterogeneity. For the "L" system, greater rail access is associated with lower ride-hailing prices in low-income communities, meaning that rail serves as a replacement and primary mode of transportation. Compared to the result in low-income groups, rail accessibility is correlated to higher ride-hailing prices in middle-income regions, likely due to concentrated demand around transit hubs. In high-income communities, the effect is minimal and slightly negative, consistent with private-car-dominant travel patterns. As I expected, both trip distance and trip span are strong positive predictors across all income groups. This is the same as the traditional taxi industry, where pricing is still based on price per mile and price per minute. In addition, it shows that high-income individuals exhibit the greatest sensitivity to trip distance.

Table 4: Estimated Coefficients on Control Variables across Income Groups

	Pooled	Low Income	Middle Income	High Income
Temperature (°C)	-0.008***	-0.003	-0.014***	-0.007***
	(0.001)	(0.003)	(0.001)	(0.001)
Wind Speed (m/s)	-0.010***	-0.015**	-0.012***	-0.009***
	(0.001)	(0.007)	(0.003)	(0.002)
Precipitation (mm)	0.001***	0.002	0.004***	0.001**
	(0.000)	(0.002)	(0.001)	(0.001)
Snowfall (mm)	0.007***	0.004***	0.006***	0.007***
	(0.000)	(0.001)	(0.000)	(0.000)
Bus	0.040***	0.533***	-0.301***	0.124***
	(0.006)	(0.069)	(0.032)	(0.008)
L Rail	-0.065***	-1.285***	0.375***	-0.030***
	(0.006)	(0.154)	(0.030)	(0.010)
Taxi Trips	0.044***	0.329***	0.032***	0.083***
	(0.001)	(0.083)	(0.001)	(0.002)
Peak Hour Trip	-0.003	-0.021	-0.011*	0.001
	(0.003)	(0.016)	(0.006)	(0.004)
Trip Distance (miles)	0.967***	0.871***	0.932***	1.000***
	(0.002)	(0.007)	(0.003)	(0.002)
Trip Duration (s)	0.004***	0.004***	0.004***	0.004***
	(0.000)	(0.000)	(0.000)	(0.000)
Tuesday	0.043***	0.059***	0.062**	0.075***
	(0.006)	(0.005)	(0.027)	(0.009)
Wednesday	0.042***	0.064***	0.060**	0.066***
	(0.007)	(0.006)	(0.028)	(0.010)
Thursday	0.054***	0.085***	0.024	0.103***
	(0.006)	(0.006)	(0.026)	(0.010)
Friday	0.020**	0.046***	0.016	0.099***
	(0.007)	(0.006)	(0.030)	(0.011)
The Loop	0.535***	0.574***	0.534***	0.591***
	(0.005)	(0.005)	(0.024)	(0.015)
Key Transportation or Tourist Zones	0.564***	0.618***	0.624***	0.797***
	(0.005)	(0.004)	(0.020)	(0.011)
Observations	1385310	53645	410354	921311

Notes: Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

I also want to briefly discuss the remaining dummy variables, including the day of the week and the after-policy area types. The main underlying finding is that the "slow start on Monday" phenomenon does exist, as evidenced by the higher asking prices on other weekdays (even if not significantly on some days). As for area type, the results in the table mainly show that the policy is indeed effective. These findings on controllers underscore the substantial heterogeneity in travel behavior and its responsiveness to environmental and contextual factors across different income levels.

Table 5: Estimated Treatment Effects of the Congestion Tax across Income Groups

	Pooled	Low Income	Mid Income	High Income
Post × TreatYear	1.931***	1.264***	1.831***	1.968***
Std. Error	(0.014)	(0.069)	(0.026)	(0.017)
Observations	1385310	53645	410354	921311

Notes: Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Having confirmed that the control variables are as expected, the analysis now turns to the core coefficients that I mentioned in the previous section. Table 5 reports the estimated treatment effects of the congestion tax across different income groups. The coefficients for the interaction term $Post \times TreatYear$ are positive and statistically significant in all subsamples, exhibiting that the congestion tax leads to a consequential increase in the price of ride-hailing. The influence is most pronounced among high-income communities, with an estimated coefficient of 1.968, followed by the pooled sample (1.931) and middle-income areas (1.831). Notably, while the effect is slighter for low-income communities (1.264), it remains statistically significant, indicating that the policy has a measurable, albeit less noticeable effect. As an addition to the description, the R^2 of the pooled regression is approximately 0.78. The R^2 values are around 0.81 for the low-income group, 0.77 for the middle-income group, and 0.80 for the high-income group. Given the result, there exists a heterogeneous effect of the congestion tax, with high-income travelers experiencing the strongest response.

This conclusion can be further confirmed by the Wald test, and the Wald tests confirm that the price increase following the congestion tax is significantly larger in high-income and middle-income areas compared to low-income areas (p < 0.01 for both). In detail, the largest divergence is marked between high-income and low-income neighborhoods (Wald = 98.14), followed by middle-income versus low-income (Wald = 59.13), indicating that we do live in a society with substantial heterogeneity in pricing dynamics across the income spectrum.

7.4 Robustness Checks

Table 6: Bandwidth Robustness Checks across Income Groups

Bandwidth	Pooled	Low Income	Mid Income	High Income
±5 days	2.869*** (0.000)	2.099*** (0.448)	2.671*** (0.165)	2.936*** (0.109)
±7 days	2.106*** (0.044)	1.065*** (0.222)	1.968*** (0.082)	2.148*** (0.054)
±10 days	1.900*** (0.018)	1.314*** (0.089)	1.802*** (0.034)	1.934*** (0.022)
±15 days	1.931*** (0.014)	1.264*** (0.069)	1.831*** (0.026)	1.968*** (0.017)
±30 days	1.967*** (0.010)	1.364*** (0.047)	1.807*** (0.018)	2.041*** (0.012)

Notes: This table reports the estimated treatment effects under different bandwidths around the policy cutoff. Each cell shows the coefficient and the robust standard error (in parentheses). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

In addition to completing the primary cause-effect identification, I also conducted bandwidth robustness checks for our Diff-in-RDiT model, as shown in Table 6 and Appendix Figure A11. To ensure the robustness of my findings, I tested bandwidths of 5, 7, 10, 15, and 30 days. Table 6 documents bandwidth robustness inspections, demonstrating that the estimated treatment effects remain positive and statistically significant across a broad range of window widths. Effect magnitudes were highly agreeing, specifically in middle-income and high-income communities. While low-income areas exhibit slightly smaller and less precise estimates, especially at narrower bandwidths. The overall pattern reinforces the bandwidth robustness of my findings.

Table 7: Treatment effects across trend orders.

Trend Order	Pooled	Low Income	Mid Income	High Income
Order 1	1.931***	1.264***	1.831***	1.968***
	(0.014)	(0.069)	(0.026)	(0.017)
Order 3	2.740***	1.962***	2.569***	2.741***
	(0.081)	(0.397)	(0.153)	(0.099)

Notes: This table reports the estimated treatment effects under different trend orders. Each cell shows the coefficient and the robust standard error (in parentheses). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

I also verified whether the results are stable between polynomial time trend functions of different orders. By setting different polynomial time functions, I obtained results with different values, but significant differences between groups and consistent patterns of group differences.

Table 8: Bandwidth Robustness Checks across Income Groups

Bandwidth	Pooled	Low Income	Mid Income	High Income
±5 days	0.243***	0.162***	0.220***	0.251***
	(0.000)	(0.033)	(0.012)	(0.008)
±7 days	0.182***	0.095***	0.166***	0.188***
	(0.003)	(0.016)	(0.006)	(0.004)
±10 days	0.166***	0.115***	0.153***	0.171***
	(0.001)	(0.006)	(0.002)	(0.002)
±15 days	0.162***	0.108***	0.150***	0.167***
	(0.001)	(0.005)	(0.002)	(0.001)
±30 days	0.165***	0.112***	0.151***	0.173***
	(0.001)	(0.003)	(0.001)	(0.001)

Notes: This table also reports the estimated treatment effects under different bandwidths around the policy cutoff. Compared with the previous regression, these results are regressed on log ride-hailing prices. Hence, they provide estimates on the change in proportion. Each cell shows the coefficient and the robust standard error (in parentheses). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Finally, I also verified whether the price adjustment heterogeneity found in the study is mechanism-stable. The same group differences and difference patterns can be obtained through the same regression by replacing the dependent variable with the log value of the ride-hailing price. Based on these new regression results, I can still conclude that after the congestion tax policy is implemented, the ride-hailing platform has the highest price increase in high-income communities and the lowest price increase in low-income communities.

8 Discussion

The empirical results of this study provide convincing evidence that ride-hailing platforms adjust prices in response to regulatory shocks in a manner that is systematically related to community income characteristics. Specifically, after the congestion tax was imposed, fares in high-income areas increased the most, while low-income communities saw significantly smaller price adjustments. Based on a study window of thirty days before and after the policy took effect and assuming a linear time trend, it can be estimated that the platform increased fares by $1.968 for high-income communities, $1.831 for middle-income communities, and $1.264 for low-income communities. This finding reveals between-group heterogeneity in platform pricing across different communities. In the current analytical model, the effects of different types of policy shocks on the origin and destination of trips have been controlled. In addition, differences in key trip characteristics and timing have been accounted for, and the impact of public transportation accessibility and taxi transactions across communities has been considered. Therefore, the current causal inference framework can be sufficiently confident to determine that the between-group differences found come from the platform's price discrimination.

This finding is of great significance for understanding the behavioral logic of pricing algorithms. This is because the finding supports the following theoretical view: platforms can use agent-based Bayesian pricing mechanisms to perform third-degree price discrimination on users, that is, platforms use observable signals (such as geographic location) as a proxy for latent attributes (such as willingness to pay). Although ride-hailing companies publicly claim that the purpose of dynamic pricing is simply to balance real-time supply and demand, our results show that pricing may also be affected by inferred neighborhood consumer characteristics. From a theoretical perspective, this is consistent with the third-degree price discrimination model, especially under conditions of partial observability and legal constraints on the use of personal data. In this case, the platform does not need to directly obtain individual income levels, but can rely on neighborhood indicators as an effective alternative to economic status analysis.

Based on classic price discrimination theory, price discrimination involves adjusting prices according to users' demand elasticity. Therefore, based on my current research results, it appears that users in high-income communities have lower demand elasticity compared to users in low-income communities. However, this conclusion may be subject to skepticism. Critics might argue that high-income groups generally have more travel options and are less dependent on ride-hailing services. While this argument holds some truth, it only captures part of the story. Jawaher Binsuwadan et al. (2023) introduced another important perspective: the value of travel time savings is also a measurable factor [Binsuwadan et al., 2023]. Their research found that as income increases, people's valuation of saving travel time also rises. This suggests that high-income individuals are more willing to pay premium prices to save time, which is consistent with the findings of my study. As a community's overall income level rises, residents' valuation of time savings through ride-hailing services increases, allowing platforms to charge higher prices. Naturally, this also implies that premium fast services (such as Uber Black) may find greater success in these areas. However, this phenomenon raises worries about whether the pursuit of efficient allocation may also be accompanied by growing fairness imbalances. A study by Sicheng Wang et al. (2024), also based on Chicago's ride-hailing data, discovered that high-income and low-minority communities had quite shorter ride-hailing travel times to fundamental services such as medical facilities, restaurants, and grocery stores [Wang et al., 2024]. In contrast, residents of low-income and high-minority communities experienced longer ride-hailing trips, indicating poorer accessibility to basic services. Ideally, ride-hailing is seen as a way to bridge mobility gaps, but in reality, it appears to exacerbate economic disparities between urban communities.

What is more concerning is that this hidden inequality has worsened following the implementation of the congestion tax. Citing Justin Tian et al.'s study (2024) on the same policy, it was found that after the congestion fee was introduced, the number of ride-hailing trips in low-income areas dropped by 11.5%, a larger decline than in high-income areas [Tian et al., 2024]. Thus, while the smaller price additions for low-income communities after the policy might seem like a friendly strategy to maintain more price-sensitive and less time-sensitive customers, it truly mirrors an "unconscious" amplification of travel inequality in the market. Notably, this widening gap is not merely a side effect of isolated policy changes but reflects deeper structural dynamics within such platform economies. As underlined by Arjan de Ruijter et al. (2024), ride-hailing platforms are more likely to thrive in environments characterized by socioeconomic inequality [de Ruijter et al., 2024]. Since inequality societies simultaneously provide a larger supply of cheap labor (faster growth on the supply side) and a greater number of passengers willing to pay higher prices to save time (growth on the demand side). In this light, the observed pricing heterogeneity not only challenges the assumption of algorithmic neutrality but also reveals how platform profit maximization strategies may reinforce and exacerbate existing social inequalities, someday undermining broader social welfare.

9 Conclusion and Future Work

Recalling my research, my study aimed to investigate whether ride-hailing platforms engage in price discrimination across communities with different socioeconomic features. I use a natural experiment, Chicago's 2020 congestion tax on ride-hailing services, to apply the Diff-in-RDiT identification framework to rich datasets. My results show significant differences in price reactions between communities under a uniform policy: high-income communities underwent the largest price increase of \$1.968, while low-income areas experienced the smallest increase of only \$1.264. These findings imply that pricing algorithms may implicitly incorporate community-level signals, such as income and poverty rates, when adjusting prices, evoking a rethinking of platform fairness and algorithmic justice.

My study also equips empirical evidence for the existence of proxy-based third-degree price discrimination on ride-hailing, contributing to the emerging literature on platform economics. Given clustering communities using socioeconomic indicators at the census tract level and combining travel, traffic, and weather data, the study delivers a powerful framework to identify heterogeneous effects in urban travel systems. Briefly, my findings demonstrate that even in the absence of explicit discrimination, opaque pricing mechanisms can exacerbate transportation inequalities.

From a policy perspective, the findings underscore the importance of transparency in algorithmic pricing. While dynamic pricing serves efficiency goals, the structural inequalities it mirrors and the social disparities it amplifies deserve closer scrutiny, specifically when private companies control public resources and citizen information.

It is worth noting that although my identification strategy and robustness tests enhance the credibility of causal inferences, there are still some limitations. For example, aggregating socioeconomic indicators at the census tract level may mask differences within communities, and the lack of individual-level user data also hinders more detailed research on personalized pricing adjustments. Even at the community level, the lack of precise data on boarding locations limits our ability to detect geographic differences between neighboring communities. If non-anonymized data are available, more powerful follow-up findings may be obtained through more detailed research under academic ethical supervision. In addition, the application of community-linked data to bus systems is still imperfect. The current method of estimating bus ridership - calculated based on the proportion of bus route length within the community to the total route length - has low accuracy. For example, a bus route may span a large census tract but have few actual stops. This point, if per-stop flow data similar to the "L" system can be obtained, will address this estimation defect. In addition, although our design controls for weather, alternative transportation options, and time fixed effects, unobservable factors such as local events or demand shocks in specific communities may still bring residual confounding effects.

A particularly important direction for future work is to examine tax pass-through and tax incidence associated with Chicago's congestion tax. While this paper documents heterogeneous fare adjustments across neighborhoods, it does not separately identify how the tax burden is divided between riders, drivers, and the platform. Incorporating a welfare-economic framework, including estimating pass-through rates, incidence heterogeneity, and distributional welfare effects, would provide a more complete understanding of the economic consequences of the policy. Given the quasi-experimental setting and the richness of the TNP dataset, embedding the current design into a formal incidence model represents a promising and natural extension.

Future research could extend this analysis by incorporating data from other urban centers to test its cross-city generalizability or by exploring interactions with additional demographic dimensions such as race, education, or smartphone usage. In addition, it would be important to examine whether platform behavior evolves in response to changes in regulatory frameworks designed to promote algorithmic transparency and anti-discrimination compliance.

Appendix

A. Supplementary Figures

A.1. Ride-Hailing Trip Characteristics by Area Type (2019)

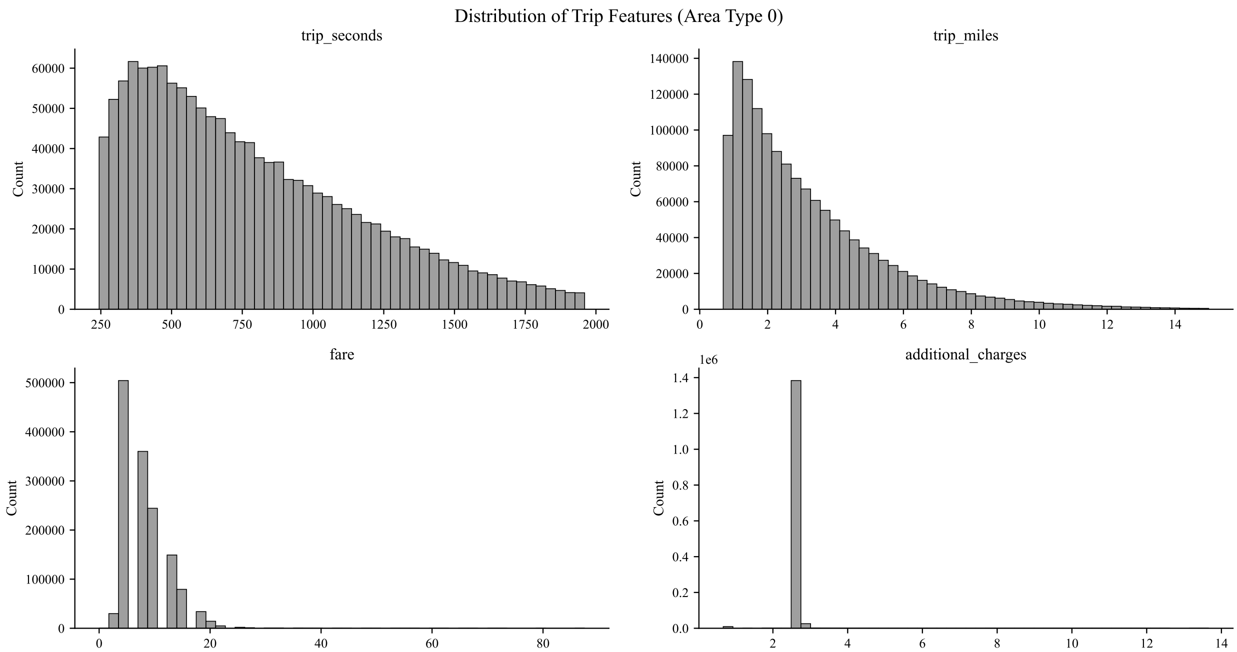

Figure A1: **Distribution of Ride-Hailing Trip Features by Area Type (2019).** Each panel shows the distribution of ride-hailing trip-level characteristics in 2019, disaggregated by area type. Area type 0 refers to general non-central neighborhoods, area type 1 corresponds to the Loop area (central business district), and area type 2 covers key transportation or tourist zones such as O'Hare Airport and Near North Side. The features include trip duration (in seconds), trip distance (in miles), total fare (in USD), and additional charges (e.g., surcharges). All distributions are shown at the trip level.

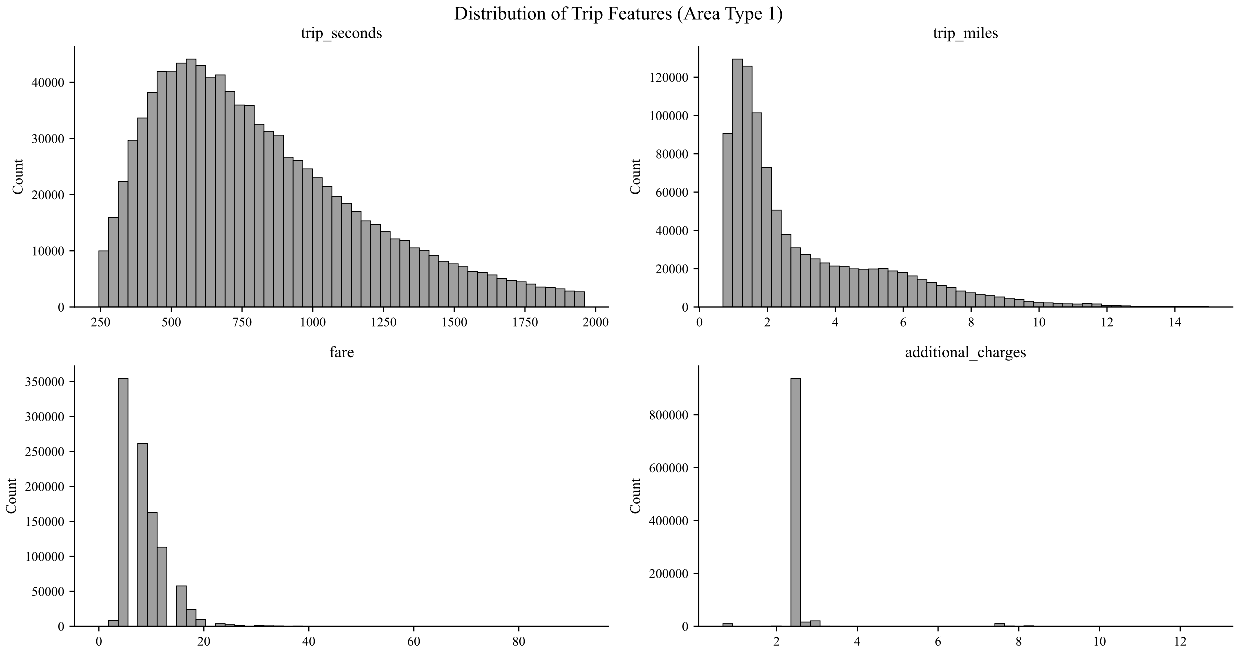

Figure A2: (Same description as above)

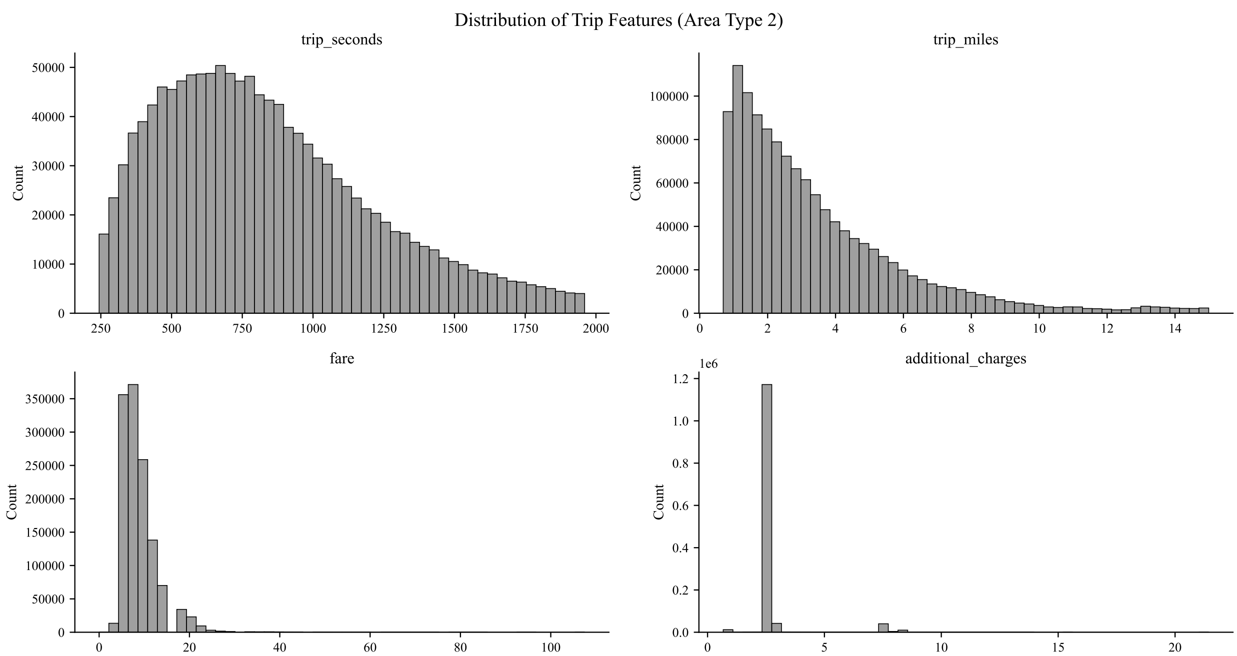

Figure A3: (Same description as above)

A.2. Ride-Hailing Trip Characteristics by Area Type (2020)

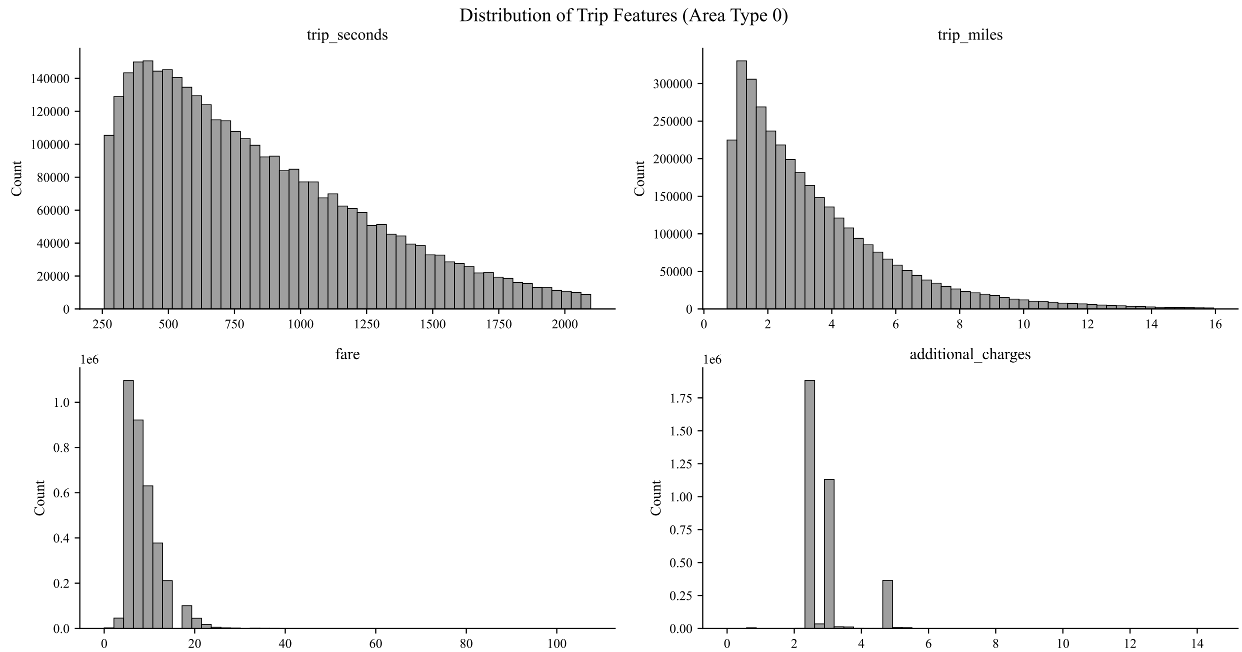

Figure A4: **Distribution of Ride-Hailing Trip Features by Area Type (2020).** Each panel shows the distribution of ride-hailing trip-level characteristics in 2020, disaggregated by area type. Area type 0 refers to general non-central neighborhoods, area type 1 corresponds to the Loop area (central business district), and area type 2 covers key transportation or tourist zones such as O'Hare Airport and Near North Side. The features include trip duration (in seconds), trip distance (in miles), total fare (in USD), and additional charges (e.g., surcharges). All distributions are shown at the trip level.

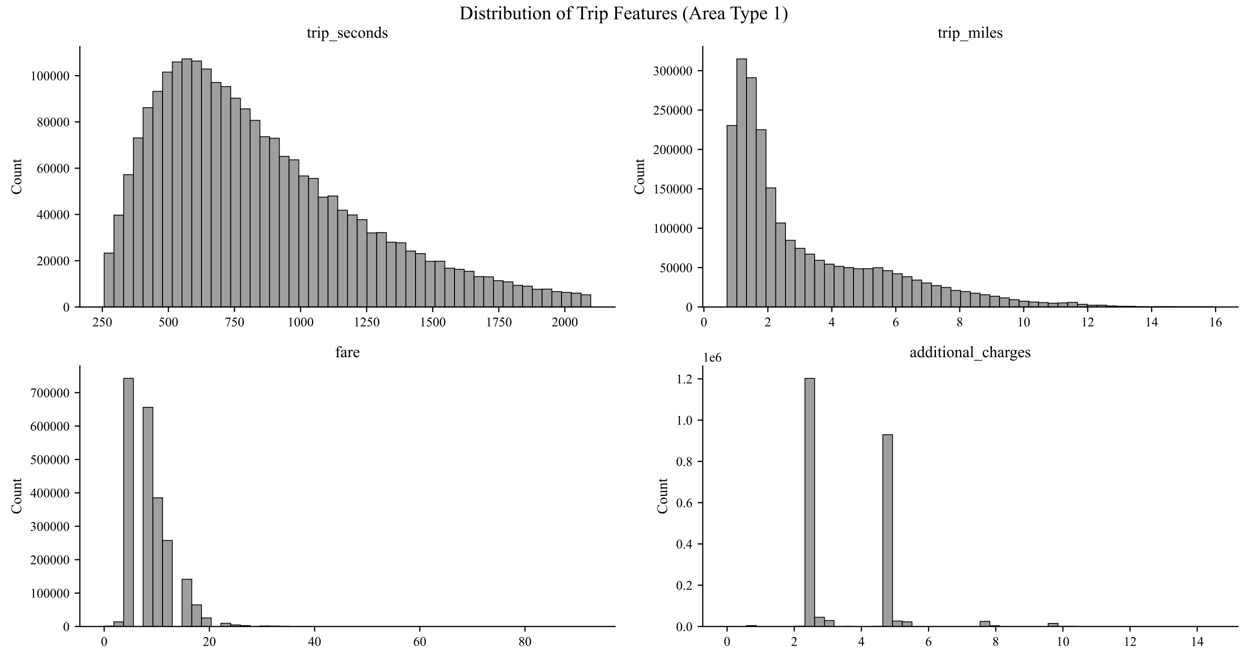

Figure A5: (Same description as above)

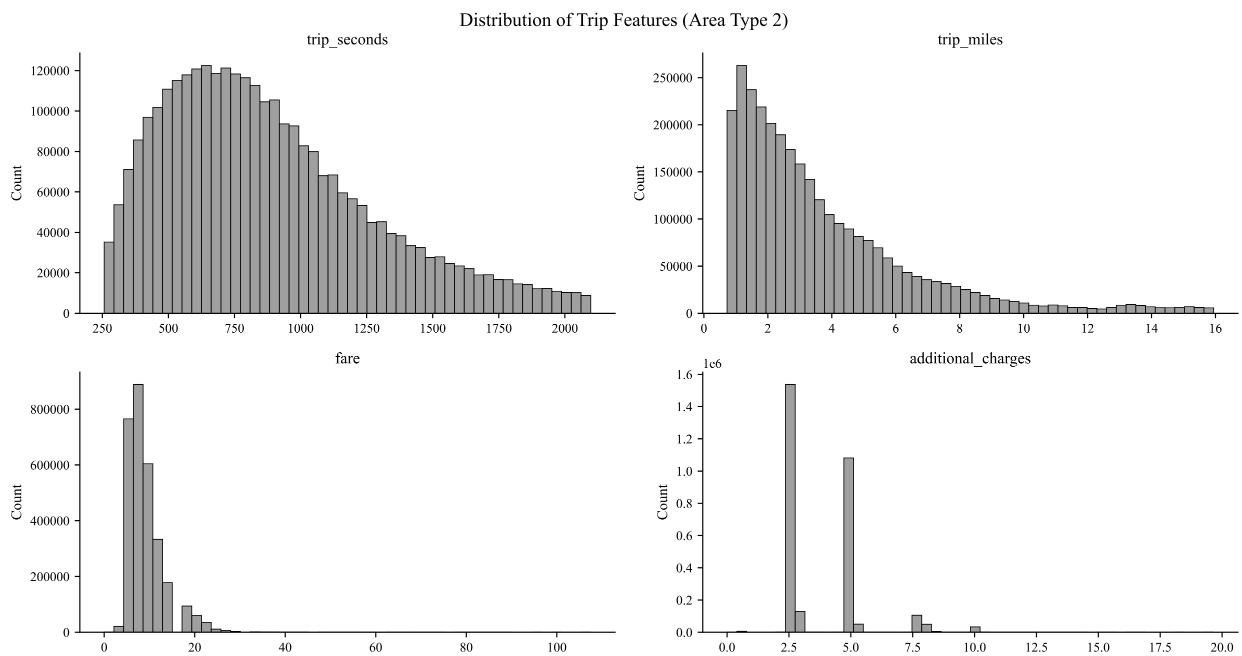

Figure A6: (Same description as above)

A.3. Temporal Distribution of Ride-Hailing Trips on Weekdays

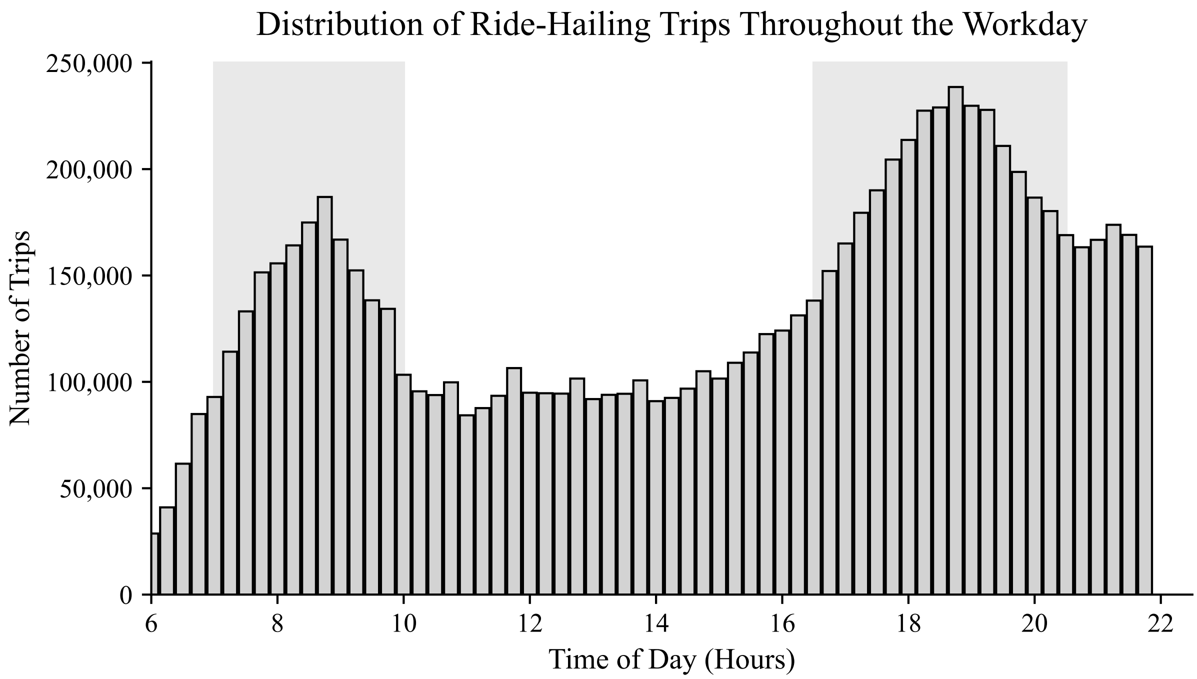

Figure A7: **Distribution of Ride-Hailing Trips Throughout the Workday.** This figure shows the distribution of ride-hailing trips across the workday, using 15-minute intervals from 6:00 AM to 10:00 PM. The number of trips sharply increases during the typical morning rush hour (7:00–10:00 AM) and again during the evening rush hour (4:30–8:30 PM), reflecting commuting and post-work travel patterns. These peaks highlight the temporal concentration of demand for ride-hailing services in Chicago.

A.4. Income and Poverty Distribution Across Census Tracts

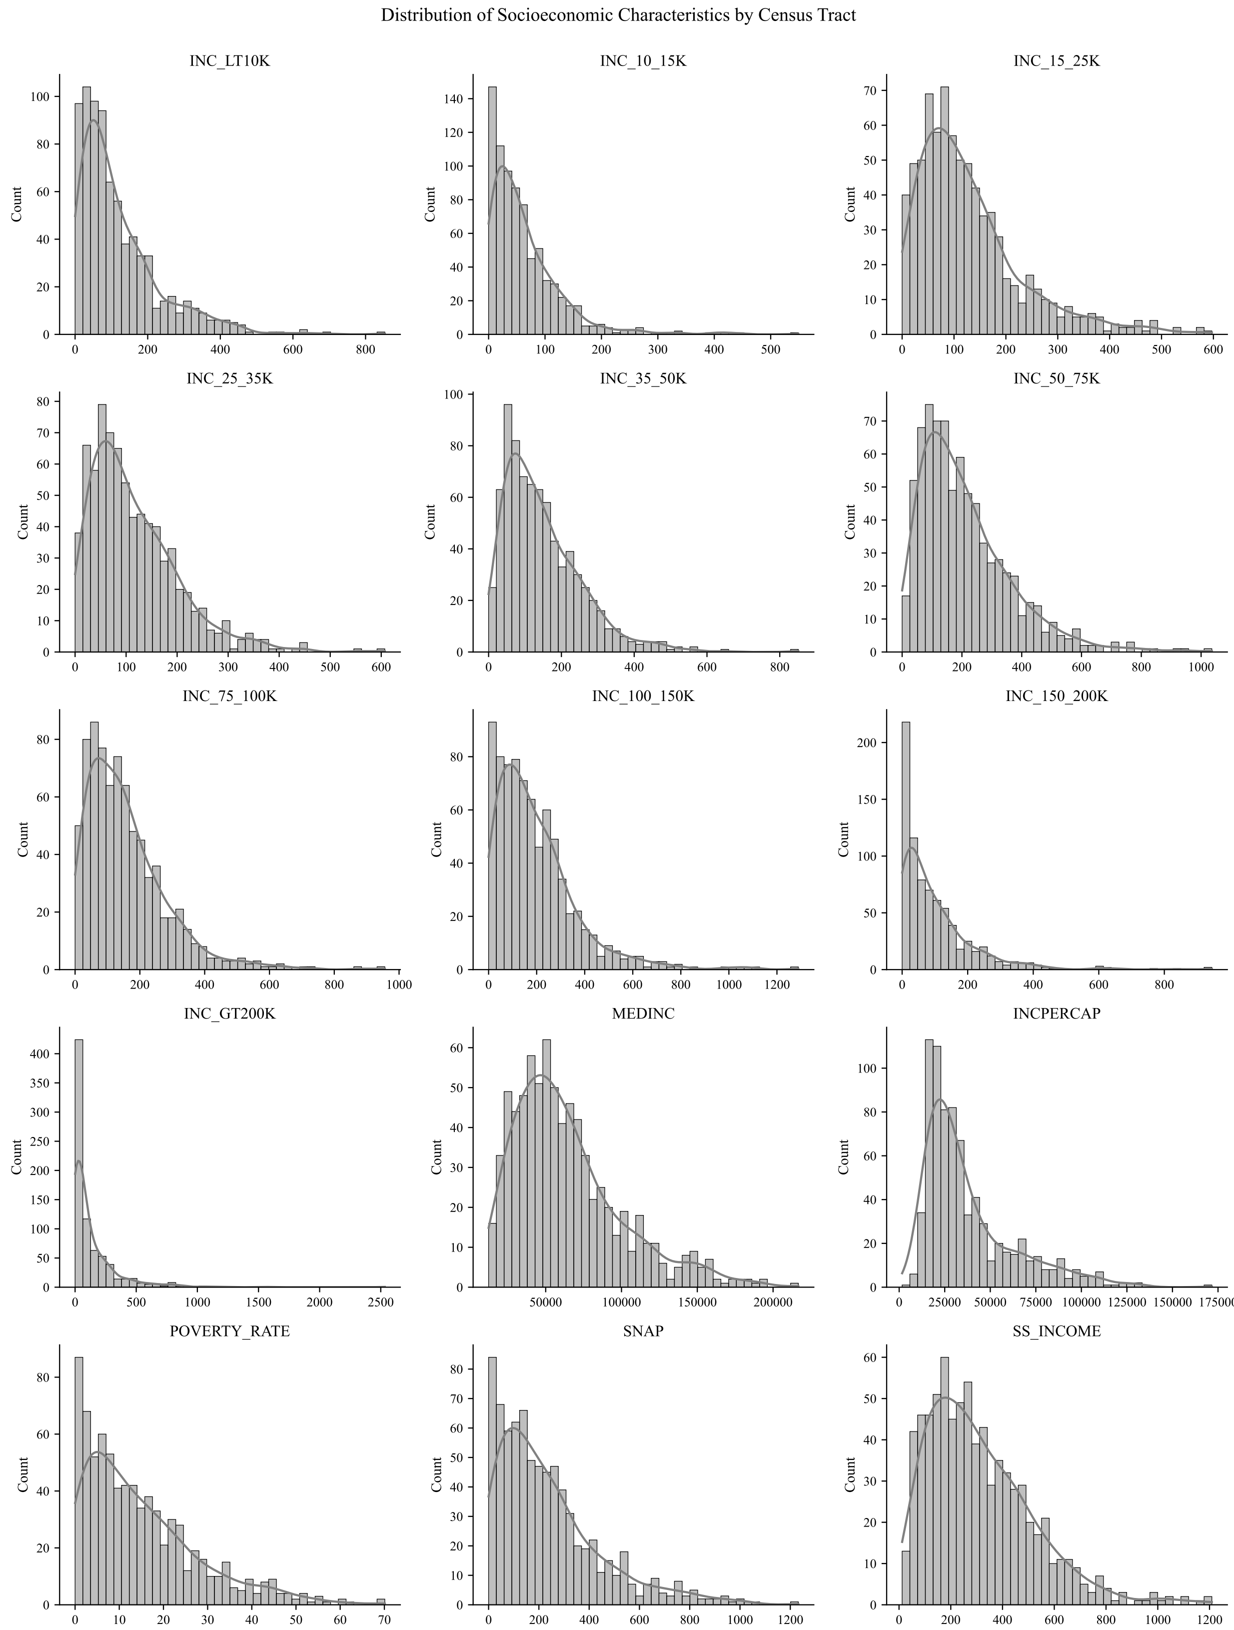

Figure A8: **Distribution of Socioeconomic Characteristics by Census Tract.** Each panel displays the distribution of a socioeconomic variable across census tracts in Chicago. The variables include median household income, per capita income, income bracket counts, poverty rate, participation in food stamp (SNAP) programs, and households with social security income. Most variables exhibit strong right-skewed distributions. This skewness motivates the use of log-transformations before applying K-means clustering and standardization, as such transformations help reduce the influence of outliers and improve the comparability of feature scales.

A.5. Spatial Mapping of Socioeconomic Indicators in Chicago

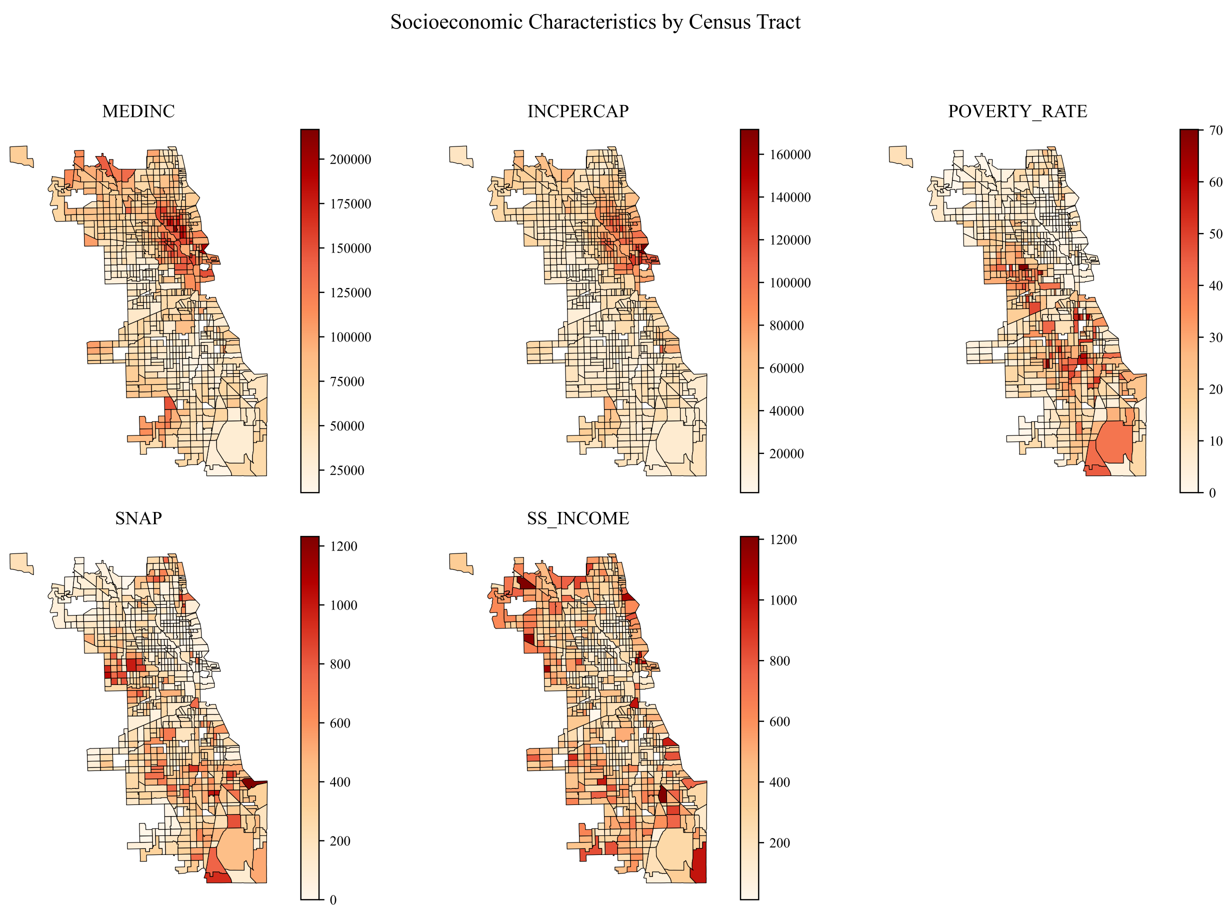

Figure A8: **Spatial Distribution of Socioeconomic Characteristics in Chicago.** Each panel shows the spatial distribution of socioeconomic characteristics at the census tract level in Chicago. Color intensity reflects the level of the variable.

Figure A9: (Same description as above)

A.6. Clustered Income Groups Based on Socioeconomic Indicators

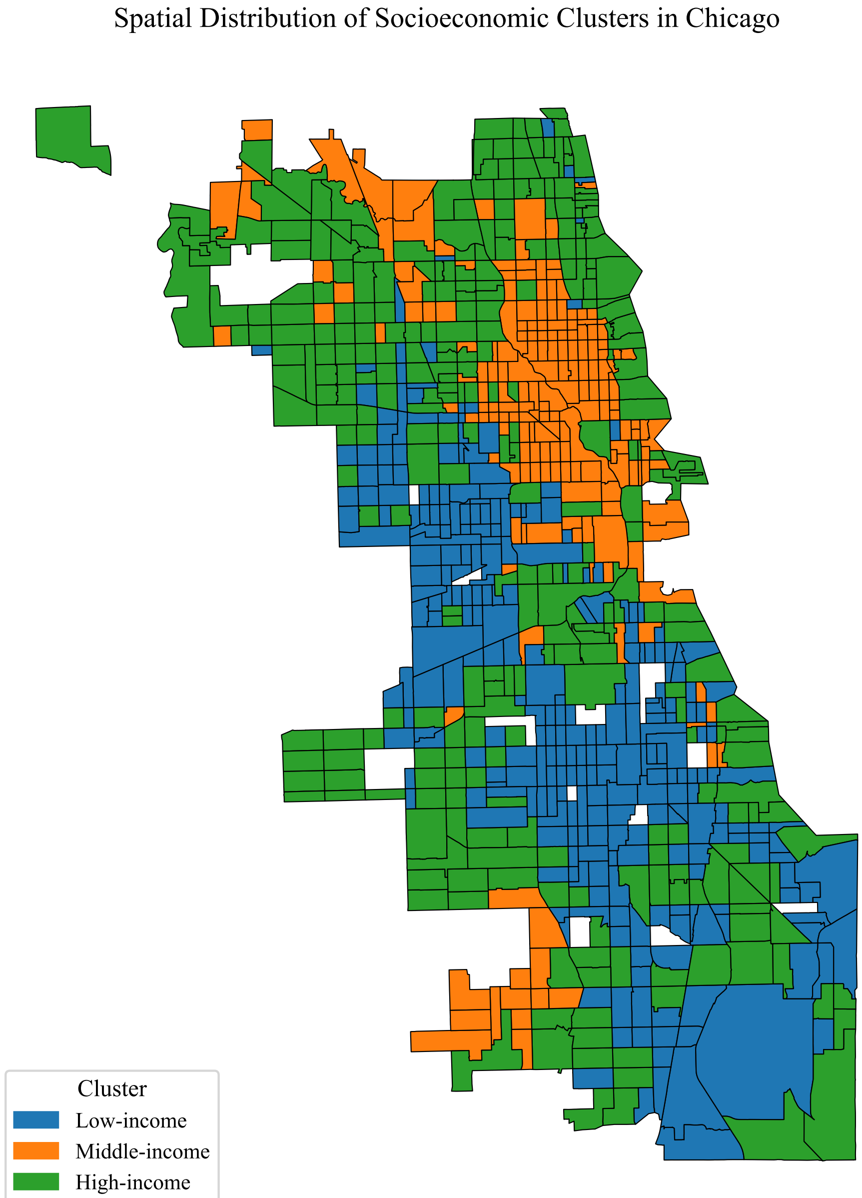

Figure A10: **Spatial Distribution of Clustered Census Tract in Chicago.** This map displays the spatial distribution of census tract-level clusters derived from K-means classification based on socioeconomic indicators. Each cluster represents a distinct neighborhood profile, characterized by income composition, poverty rates, and reliance on social support.

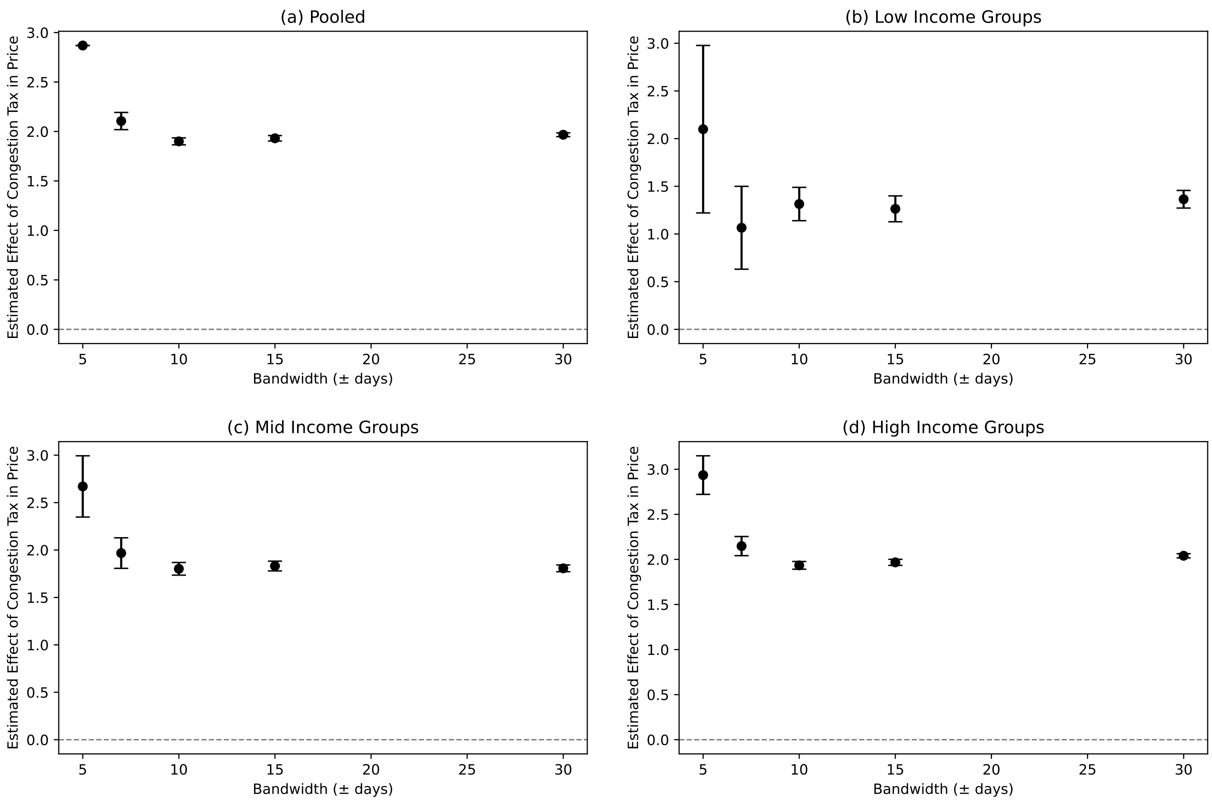

Figure A11: **Estimated Treatment Effects Across Bandwidths** This plot shows how estimated treatment effects vary across different bandwidths for the different samples, providing visual insight into the robustness of the main results. Subplot (b) represents bandwidth sensitivity for neighborhoods in the lowest socioeconomic cluster. Subplot (c) represents bandwidth sensitivity for mid-income areas. Subplot (d) represents robustness results for affluent neighborhoods.

A.7. Robustness Checks: Treatment Effects Across Bandwidths

B. Supplementary Tables

B.1. Variance Inflation Factors (VIFs) for Control Variables

Table A1: Variance Inflation Factors for Control Variables on Treatment Year

Variable	VIF
const	13.61
"L" Rail System Rides	13.60
Bus Rides	13.40
Trip Distance (miles)	2.41
Trip Duration (s)	2.37
Snowfall (mm)	1.30
Precipitation (mm)	1.28
Taxi Trips	1.27
Temperature (°C)	1.24
Peak Hour Trip	1.05
Wind Speed (m/s)	1.03

Note: This table reports variance inflation factors (VIF) for control variables in the placebo year regression. VIF values above 5 are commonly interpreted as indicating potential multicollinearity.

Table A2: Variance Inflation Factors for Control Variables on Placebo Year

Variable	VIF
const	16.11
Bus Rides	14.00
"L" Rail System Rides	13.98
Trip Duration (s)	2.31
Trip Distance (miles)	2.27
Snowfall (mm)	1.44
Precipitation (mm)	1.43
Temperature (°C)	1.28
Wind Speed (m/s)	1.11
Peak Hour Trip	1.06
Taxi Trips	1.00

Note: This table reports variance inflation factors (VIF) for control variables in the placebo year regression. VIF values above 5 are commonly interpreted as indicating potential multicollinearity.

B.2. PCA Dimension Reduction

Table A3: PCA Loadings on Socioeconomic Variables

	PC1	PC2
Median Household Income	-0.398	-0.015
Per Capita Income	-0.372	0.001
Income $<$ 10k	0.208	-0.206
Income 10k–15k	0.198	-0.176
Income 15k–25k	0.229	-0.281
Income 25k–35k	0.150	-0.335
Income 35k–50k	0.078	-0.367
Income 50k–75k	-0.063	-0.406
Income 75k–100k	-0.175	-0.345
Income 100k–150k	-0.263	-0.300
Income 150k–200k	-0.308	-0.202
Income $>$ 200k	-0.335	-0.113
Poverty Rate (%)	0.357	0.014
Households on SNAP	0.317	-0.173
Households with Social Security	0.001	-0.378

Notes: This table reports the loadings of the first two principal components extracted from standardized socioeconomic variables. Loadings represent the contribution of each original variable to the corresponding principal component.

References

Niels Agatz, Alan Erera, Martin Savelsbergh, and Xing Wang. Optimization for dynamic ride-sharing: A review. *European Journal of Operational Research*, 223(2):295–303, 2012. ISSN 0377-2217. doi: https://doi.org/10.1016/j.ejor.2012.05.028. URL https://www.sciencedirect.com/science/article/pii/S0377221712003864.

Kay Axhausen, Andrea Zimmermann, Stefan Schönfelder, Guido Rindsfüser, and Thomas Haupt. Observing the rhythms of daily life: A six-week travel diary. *Transportation*, 29:95–124, 05 2002. doi: 10.1023/A:1014247822322.

Siddhartha Banerjee, Ramesh Johari, and Carlos Riquelme. Dynamic pricing in ridesharing platforms. *SIGecom Exch.*, 15(1):65–70, September 2016. doi: 10.1145/2994501.2994505. URL https://doi.org/10.1145/2994501.2994505.

John Manuel Barrios, Yael V. Hochberg, and Hanyi Yi. The cost of convenience: Ridehailing and traffic fatalities. SSRN Working Paper, April 3 2019. Available at SSRN: https://ssrn.com/abstract=3361227 or http://dx.doi.org/10.2139/ssrn.3361227.

Patrick Bayer, Fernando Ferreira, and Robert McMillan. A unified framework for measuring preferences for schools and neighborhoods. *Journal of Political Economy*, 115(4):588–638, 2007. ISSN 00223808, 1537534X. URL http://www.jstor.org/stable/10.1086/522381.

Jawaher Binsuwadan, Mark Wardman, Gerard de Jong, Richard Batley, and Phill Wheat. The income elasticity of the value of travel time savings: A meta-analysis. *Transport Policy*, 136:126–136, 2023. ISSN 0967-070X. doi: https://doi.org/10.1016/j.tranpol.2023.03.013. URL https://www.sciencedirect.com/science/article/pii/S0967070X23000720.

Abel Brodeur and Kerry Nield. An empirical analysis of taxi, lyft and uber rides: Evidence from weather shocks in nyc. *Journal of Economic Behavior & Organization*, 152:1–16, 2018. ISSN 0167-2681. doi: https://doi.org/10.1016/j.jebo.2018.06.004. URL https://www.sciencedirect.com/science/article/pii/S0167268118301598.

J. Massey Cashore, Peter I. Frazier, and Eva Tardos. Dynamic pricing provides robust equilibria in stochastic ridesharing networks, 2022. URL https://arxiv.org/abs/2205.09679.

Juan Camilo Castillo, Daniel T. Knoepfle, and Eric Glen Weyl. Matching and pricing in ride hailing: Wild goose chases and how to solve them. https://ssrn.com/abstract=2890666, February 2024. Available at SSRN: https://ssrn.com/abstract=2890666 or http://dx.doi.org/10.2139/ssrn.2890666.

Yenjae Chang, Clifford Winston, and Jia Yan. Does uber benefit travelers by price discrimination? *Journal of Law and Economics*, 65(S2):S433 – S459, 2022. URL https://EconPapers.repec.org/RePEc:ucp:jlawec:doi:10.1086/721266.

M. Keith Chen. Dynamic pricing in a labor market: Surge pricing and flexible work on the uber platform. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, page 455, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450339360. doi: 10.1145/2940716.2940798. URL https://doi.org/10.1145/2940716.2940798.

Raj Chetty, Nathaniel Hendren, and Lawrence F. Katz. The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment. *American Economic Review*, 106(4):855–902, April 2016. doi: 10.1257/aer.20150572. URL https://www.aeaweb.org/articles?id=10.1257/aer.20150572.

Regina R. Clewlow and Gouri Shankar Mishra. Disruptive transportation: The adoption, utilization, and impacts of ride-hailing in the united states. *UC Davis Sustainable Transportation Energy Pathways of the Institute of Transportation Studies*, 2017. URL https://itspubs.ucdavis.edu/publication_detail.php?id=2750.

Judd Cramer and Alan B Krueger. Disruptive change in the taxi business: The case of uber. Working Paper 22083, National Bureau of Economic Research, March 2016. URL http://www.nber.org/papers/w22083.

Janet Currie and Aaron Yelowitz. Are public housing projects good for kids? *Journal of Public Economics*, 75(1):99–124, 2000. ISSN 0047-2727. doi: https://doi.org/10.1016/S0047-2727(99)00065-1. URL https://www.sciencedirect.com/science/article/pii/S0047272799000651.

Arjan de Ruijter, Oded Cats, and Hans van Lint. Ridesourcing platforms thrive on socio-economic inequality. *Scientific Reports*, 14(1):7371, 2024. ISSN 2045-2322. doi: 10.1038/s41598-024-57540-x. URL https://doi.org/10.1038/s41598-024-57540-x.

Manasi Deshpande. Does welfare inhibit success? the long-term effects of removing low-income youth from the disability rolls. *American Economic Review*, 106(11):3300–3330, November 2016. doi: 10.1257/aer.20151129. URL https://www.aeaweb.org/articles?id=10.1257/aer.20151129.

Yanbo Ge, Christopher R. Knittel, Don MacKenzie, and Stephen Zoepf. Racial and gender discrimination in transportation network companies. *NBER Working Paper No. w22776*, October 2016. doi: 10.2139/ssrn.2861708. URL https://ssrn.com/abstract=2861708.

Jonathan D. Hall, Craig Palsson, and Joseph Price. Is uber a substitute or complement for public transit? *Journal of Urban Economics*, 108:36–50, 2018. ISSN 0094-1190. doi: https://doi.org/10.1016/j.jue.2018.09.003. URL https://www.sciencedirect.com/science/article/pii/S0094119018300731.

Benjamin Hansen. Punishment and deterrence: Evidence from drunk driving. *American Economic Review*, 105(4):1581–1617, April 2015. doi: 10.1257/aer.20130189. URL https://www.aeaweb.org/articles?id=10.1257/aer.20130189.

Alejandro Henao and Wesley E. Marshall. The impact of ride-hailing on vehicle miles traveled. *Transportation*, 46(6):2173–2194, December 2019. doi: 10.1007/s11116-018-9923-2. URL https://ideas.repec.org/a/kap/transp/v46y2019i6d10.1007_s11116-018-9923-2.html.

James O. Huff and Susan Hanson. Repetition and variability in urban travel. *Geographical Analysis*, 18(2): 97–114, 1986. doi: https://doi.org/10.1111/j.1538-4632.1986.tb00085.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1538-4632.1986.tb00085.x.

Sarah Miller, Norman Johnson, and Laura R Wherry. Medicaid and mortality: New evidence from linked survey and administrative data*. *The Quarterly Journal of Economics*, 136(3):1783–1829, 01 2021. ISSN 0033-5533. doi: 10.1093/qje/qjab004. URL https://doi.org/10.1093/qje/qjab004.

Farid M. Naini, Jayakrishnan Unnikrishnan, Patrick Thiran, and Martin Vetterli. Where you are is who you are: User identification by matching statistics. *Trans. Info. For. Sec.*, 11(2):358–372, February 2016. ISSN 1556-6013. doi: 10.1109/TIFS.2015.2498131. URL https://doi.org/10.1109/TIFS.2015.2498131.

Akshat Pandey and Aylin Caliskan. Disparate impact of artificial intelligence bias in ridehailing economy's price discrimination algorithms. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '21, page 822–833, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450384735. doi: 10.1145/3461702.3462561. URL https://doi.org/10.1145/3461702.3462561.

Lisa Rayle, Danielle Dai, Nelson Chan, Robert Cervero, and Susan Shaheen. Just a better taxi? a survey-based comparison of taxis, transit, and ridesourcing services in san francisco. *Transport Policy*, 45:168–178, 2016. ISSN 0967-070X. doi: https://doi.org/10.1016/j.tranpol.2015.10.004. URL https://www.sciencedirect.com/science/article/pii/S0967070X15300627.

Janet L. Smith, Zafer Sonmez, and Nicholas Zettel. *Growing Income Inequality and Socioeconomic Segregation in the Chicago Region*, pages 349–369. Springer International Publishing, Cham, 2021. ISBN 978-3-030-64569-4. doi: 10.1007/978-3-030-64569-4_18. URL https://doi.org/10.1007/978-3-030-64569-4_18.

Justin Tian, Min (Michelle) Chen, and Liangfei Qiu. How does ride-hailing congestion pricing shape digital access and transportation equity? evidence from a natural experiment in chicago, October 2024. URL https://ssrn.com/abstract=4975730. Available at SSRN.

Jean Tirole. *The Theory of Industrial Organization*, volume 1 of *MIT Press Books*. The MIT Press, December 1988. ISBN ARRAY(0x7e887978). URL https://ideas.repec.org/b/mtp/titles/0262200716.html.

Natasha Vajravelu. Public transit's effect on ride-hail and taxi usage. https://ssrn.com/abstract=3258592, January 2019. Available at SSRN: https://ssrn.com/abstract=3258592 or http://dx.doi.org/10.2139/ssrn.3258592.

Sicheng Wang, Robert B. Noland, and Xiao Huang. Disparities in ridehailing travel times for accessing non-work destinations. *Transportation Research Interdisciplinary Perspectives*, 28:101258, 2024. ISSN 2590-1982. doi: https://doi.org/10.1016/j.trip.2024.101258. URL https://www.sciencedirect.com/science/article/pii/S2590198224002446.

Chiwei Yan, Helin Zhu, Nikita Korolko, and Dawn Woodard. Dynamic pricing and matching in ride-hailing platforms. *Naval Research Logistics*, October 1 2018. Forthcoming.