AI-Automated Law Enforcement: An Application to Parking

Guangyu Cao Ginger Zhe Jin Jun Yan Li-An Zhou *

October 11, 2024

Abstract

This paper documents how law enforcement powered and automated by artificial intelligence (AI) affects illegal and legal parking. Drawing on data from a leading smart parking company in China, we leverage the staggered adoption of AI technology by local traffic police in a big city. Empirical results indicate that AI-automated enforcement reduces illegal parking and leads to more efficient use of parking lots, but these effects are unevenly distributed among different types of drivers.

Keywords: Artificial Intelligence (AI), Law Enforcement, Police, Parking

JEL Classification: K42, O18, R40, H00

*Cao: School of Economics, Peking University (Email: cgy1117@pku.edu.cn); Jin: University of Maryland & NBER (Email: jin@econ.umd.edu); Yan: HKU Business School, the University of Hong Kong & Guanghua School of Management, Peking University (Email: u3554522@connect.hku.hk); Zhou: Guanghua School of Management, Peking University (Email: zhoula@gsm.pku.edu.cn). Yan and Zhou have been on the board of the anonymous company, which facilitates our access to the company's proprietary data. However, the company has de-identified individual car or driver information before sharing the data access with the research team. The data analysis is required to be performed on a designated remote server with strict access control. The company does not provide any financial support to this paper, nor does it impose any restrictions on the scientific findings of the paper. We appreciate helpful comments from seminar participants at Zhejiang University. All errors are ours. Cao and Zhou acknowledge the support of the National Natural Science Foundation of China (Grant No.72473003, 72192844). The authors are listed in alphabetical order and contribute equally to the article.

1 Introduction

Many governments around the world are eager to adopt digital technology for law enforcement. But hi-tech policing – be it camera monitoring, predictive algorithms, or other artificial intelligence (AI) tools – often generates intensive debates among citizens and policymakers. Common questions include when and how the technology is used, how effective the technology is, whether the technology generates any bias or distributional effects, and what implications it has on judicial fairness, citizen privacy, and even public trust in police.

So far, researchers have found mixed evidence in a few law enforcement areas. Algorithms that aim to aid judiciary decisions (such as bailing) have the potential to improve human decision-making and reduce bias, but sometimes induce new discrimination (Kleinberg et al., 2018a; Kleinberg et al., 2018b; Ludwig & Mullainathan, 2021; Nissan, 2017). When algorithms are used to predict the geographic distribution of crimes and guide police presence, they are found to improve police productivity and decrease serious crimes (Jabri, 2021; Mastrobuoni, 2020). In the meantime, predictive policing has disproportionate racial impacts across neighborhoods (Jabri, 2021), and its implementation varies even within the same police department (Kapustin et al., 2022). In comparison, speed and red light cameras are widely used for policing in transportation. According to the US Department of Transportation (see Section 2.1 of Chapter 3 in Venkatraman et al. (2021)), camera-powered traffic enforcement, often automated, has significantly reduced injury crashes and speeds exceeding the target speed. In spite of these benefits, it has received mixed public support (more positive for red light cameras than for speed cameras) and encountered opposition out of the concerns that the program may intrude privacy, inappropriately extend law enforcement authority, and become a revenue generator rather than safety improvement.

In this paper, we document the application of AI-automated enforcement in a new area: automobile parking. While illegal parking is less severe than many other law violations, it is a basic function of local police departments, applies to a large proportion of the population on a daily basis, and affects local police departments' human resource allocation, fiscal budgets and public trust. Beyond these immediate implications for the police, parking management often has a significant impact on the whole society. In many cities, automobile is a basic – if not dominant – means of transportation, but cars are parked 95% of the time[1] and every drive has to conclude with parking. Low-speed cruising for parking positions is one of the key drivers of traffic congestion and exhaust emissions. As a result, scholars gradually realize that parking should be considered as an important "intermediate good" in the modern economy rather than a simple fixed cost of transportation (Hasker & Inci, 2014). However, most studies about parking assume parking policies are perfectly enforced and thus leave the issue of enforcement unexamined (Inci, 2015). In the meantime, limited enforcement resources, especially in large cities, make it impossible for human patrols to detect all illegal parking. The lack of accurate data has also blocked a close scrutiny of parking enforcement and its impact on society.

[1] Fortune (2016/3/13): "Today's Cars Are Parked 95% of the Time" (available at https://fortune.com/2016/03/13/cars-parked-95-percent-of-time/).

Fortunately, this problem has been addressed by some recent technological developments in AI-powered image recognition. A private Chinese company that provided data access for this article (referred to as *Company W* hereafter) specializes in AI-powered smart parking. After the traffic police procures its service, Company W installs roadside cameras to monitor curbside parking positions continuously. When drivers drive in and out, real-time pictures are analyzed by AI algorithms to detect illegal parking behavior such as double parking or parking outside the designated position. The electronic record of illegal parking is then shared with traffic police and adopted as evidence for issuing digital tickets to the registered owner of the violating car. Once the police adopts the technology, enforcement is automated, although the ticket recipient is given an opportunity to appeal and the appeal process may involve manual work.

Based on illegal parking records provided by Company W, we evaluate the effects of AI-automated enforcement on both illegal and legal parking. The data covers 968 curbside parking lots in two districts of *City Z*, which is one of the biggest cities in China and has a population over 20 million. Our empirical identification comes from the local police's staggered adoption of the technology. After Company W installed monitoring cameras and started to transmit illegal parking records to traffic police, the police did not immediately adopt the related images as evidence for ticketing. As local police officers gained more confidence in the accuracy of the AI technology, they began to adopt it in a staggered fashion, giving us an opportunity to construct both Difference-in-Differences (DID) and Regression Discontinuity Design (RDD) models to evaluate the treatment effect of adoption.

We aggregate the original illegal parking records at the parking-lot-day level. Baseline results show that AI-automated enforcement has effectively deterred illegal parking and increased drivers' compliance with parking rules. More specifically, the number of illegal parking episodes has been reduced by 6% (estimated from DID) or 3.8% (estimated from RDD), both representing a significant improvement from traditional manual enforcement. These estimates are robust to various DID and RDD specifications.

In addition, detailed data from Company W allows us to explore two dimensions that are understudied in the existing research of AI-powered law enforcement: one is the ripple effects on related legal behavior, and the other is the distributional effects across different drivers subject to the enforcement regime.

By definition, improved detection and deterrence of illegal behavior would promote lawful actions, allowing society to reap greater benefits from established rules. However, many of these ripple effects are difficult to capture – for example, better policing of serious crimes may make citizens feel safer and more willing to go out for personal pleasure, and fewer road crashes may make traffic flow smoother and more predictable. These ripple benefits are important for the benefit-cost analysis of enforcement programs, but quantifying them requires detailed data different from the records of law violations. Fortunately, the technology of Company W is used for not only detecting illegal parking but also for collecting fees for legal parking in the same parking lots. Analysis of legal parking records suggests that AI-automated enforcement has improved the usage efficiency of these parking lots. In particular, for every

parking-position-day in our sample, the number of unique parking episodes has increased by 0.267 (or 4.84%) from the sample mean before the adoption of AI-automated enforcement.

As for the distributional effects of enforcement, we find them uneven across four cuts of driver types. First, upon our request, Company W has utilized its AI algorithm to infer car brands and models based on a third-party database of car shapes. We then collect the list price of car models and decompose parking-lot-day illegal parking incidents into those by *expensive* or *cheap* cars separately. Empirical results show that drivers of cheap cars are more sensitive to AI-automated enforcement. In the other three cuts, we group parking-lot-day illegal parking incidents by (i) whether it occurs during peak or non-peak hours, (ii) whether the license plate is issued by City Z or somewhere else, and (iii) within local plates, whether the car is used for regular commute or not (based on the car's parking frequency in our data of legal parking). Results suggest that AI-automated enforcement is less effective for peak hours, for out-of-town drivers, and for regular commuters within local drivers. All these distributional effects can be explained by economic incentives: if we interpret AI-automated enforcement as an increased cost of illegal parking, one would expect people of higher economic status, people who place more value on time and location of parking, people who are either unaware of the enforcement or more prone to evade paying the ticket, and people whose costs have already approximated zero to be less price elastic when facing such a sudden increase of "price."

Our paper contributes to several strands of literature. On technology adoption in law enforcement, not only do we look at its application in a new enforcement area, the new area but also sheds light on new roles of technology. For instance, AI-automated enforcement in parking has fully replaced police patrols instead of merely being an assistant for human judges or a spatial guide for human patrols. It also provides 24/7 surveillance in the adopted areas, which is much more universal than random patrols. It gets around the classical problem of resource allocation within the police.[2] To the best of our knowledge, we are among the first to systematically demonstrate the ripple effects on legal behavior, which is an important input for the benefit-cost analysis of AI adoption. Our finding that different drivers respond differently to AI-automated enforcement is consistent with the opinion that the introduction of AI does not affect all population groups equally, either in the specific context of predictive policing (Jabri, 2021) or in a general sense (e.g., Acemoglu, 2021).

As parking is part of transportation, our paper relates to the literature on regulations in transportation. Many recent studies in this area are centered around setting speed limits or installing speed/red light cameras, which aim to enhance road safety at the cost of longer travel time (Bauernschuster & Rekers, 2022; Dušek & Traxler, 2022; Traxler et al., 2018; Venkatraman et al., 2021). As detailed below, the technology Company W uses in smart parking is more advanced (in image recognition and algorithmic judgment) than speed and red light cameras, and thus may better represent recent and future directions of AI use in transportation.

Parking, in particular, has long been ignored by economists and has typically been treated

[2]See Polinsky and Shavell (2000) and Chalfin and McCrary (2017) for review of theoretical and empirical works respectively. Hernández and Neeman (2022) have examined how to use "Bayesian persuasion" to improve deterrence of unwanted behavior such as illegal parking. Its focus on illegal parking is close to our research, but the work is purely theoretical and conducted under the framework of allocating enforcement resources.

as a fixed cost added to a trip, as summarized by Arnott (2011) and Inci (2015). This passive and neglectful view did not change until congestion and car emissions emerged as serious urban diseases in the 1990s. Since then, studies on parking have paid most attention to the trade-off between curbside parking lots and garages (Arnott & Inci, 2006; Kobus et al., 2013; Lin & Wang, 2015; Shoup, 2006) and to the pricing of parking (see Inci (2015) for a review). Recently, many cities around the world have launched "smart parking" programs involving economic measures such as dynamic pricing. For example, Krishnamurthy and Ngo (2020) find that the SFpark in San Francisco, which is composed of demand-responsive pricing and other non-price changes, has effectively increased bus ridership and reduced traffic flow. Zakharenko (2019) presents a theoretical analysis of installing sensor devices to monitor parking occupancy. However, neither touches on the role of AI or the enforcement against illegal parking. While AI-automated enforcement increases the costs of illegal parking, it does not increase the price of legal parking. In fact, it may have decreased the economic costs of legal parking if such costs include the time and stress involved in predicting parking availability and finding a parking spot when needed. Our finding on parking lot usage highlights a new efficiency gain from better parking enforcement without any explicit change in parking price.

The rest of this paper is organized as follows. Section 2 introduces the AI-automated regulation of illegal parking and its implementation in City Z. Section 3 presents the data and econometric framework. Section 4 reports baseline empirical results as well as robustness checks. Section 5 conducts additional analyses of the ripple effects on parking efficiency and distributional effects among different drivers. Section 6 concludes. To save space, we have relegated some content into three appendices. Online Appendix A provides more details about AI-automated law enforcement, Online Appendix B is dedicated for robustness checks and supplementary results of the DID model, and Online Appendix C describes the RDD strategy and reports corresponding results.

2 Research Background

2.1 AI-Automated Detection of Illegal Parking

Company W is a private Chinese company specializing in AI-powered smart parking, with an emphasis on managing parking facilities and detecting illegal parking. Traditionally, the enforcement against illegal parking was labor-intensive. The department responsible for traffic management, typically the local traffic police, employed traffic wardens or auxiliary police officers to patrol different road segments. Upon discovering illegally parked vehicles, they would issue a ticket (either handwritten or printed using handheld devices) and affix it to the windshield wiper. Given limited manpower, it is almost impossible to effectively monitor all parking spaces simultaneously. Additionally, detection by patrolling agents can be influenced by human bias, potentially leading to discriminatory enforcement.

AI-automated law enforcement has significantly transformed the enforcement against illegal parking. After the traffic police procured its services, Company W installed roadside cameras to continuously monitor curbside parking positions. When a driver parks and departs,

Company W's AI image recognition algorithm analyzes real-time images to detect potential illegal parking behavior. Once identified, the original photos and electronic records of illegal parking are shared with the traffic police and used as evidence to issue digital tickets to the registered owner of the offending vehicle.

Figure A1 provides a visual representation of the appearance of the new system. The monitoring camera is typically mounted atop a post, often a lamp post. The monitored area depicted in this image encompasses ten parking positions, spanning a total length of 50 meters. In practice, the number of parking positions a camera covers may vary, depending on the local geography and the local areas the city has designated for roadside parking. The AI-powered image recognition is conducted locally on distributed devices utilizing edge computing.[3]

In a flowchart, Figure A2 describes the process of AI-driven enforcement against illegal parking. Panel A captures a snapshot of Company W's dashboard, showcasing an illegal parking incident with specific details redacted for privacy. Once Company W forwards this record to the traffic police, the car owner receives the photographic evidence in Panel A via a mobile app and a separate text message in Panel B on their mobile phone. Upon paying the fine through the app, the driver is then issued the electronic ticket (which also serves as a payment receipt) seen in Panel C. Through this system, the enforcement of parking regulations is fully automated, enabling the traffic police to replace human patrols and manual office work in detection, ticketing, and fine collection.

The introduction of AI-automated detection can be divided into two stages. In the first stage, Company W installed monitoring cameras and began transmitting detected illegal parking incidents to the traffic police. However, during this stage, the police only reviewed the records without using them as evidence for issuing tickets. Essentially, the traffic police acted as if AI-automated detection did not exist. In the second stage, once local police officers had developed sufficient confidence in the accuracy of the AI technology, they started to utilize the AI-detected records to issue parking tickets. This was rolled out in different parking lots in a phased manner. In other words, if we consider the local police's adoption of AI-automated enforcement as the "treatment," the two-stage roll-out provides a unique opportunity to apply a Difference-in-Differences (DID) model, treating the two stages as pre- and post-treatment periods, respectively.

2.2 The Implementation in City Z

Many Chinese city governments procured services from Company W, with *City Z* being the first and the largest client. City Z is also among the largest cities in China with a population over 20 million. The data sample utilized in this paper was drawn from two districts within that city: *District X* and *District Y*.

Interestingly, prior to adopting AI-automated enforcement against illegal parking, District X had already procured services from Company W in 2018 for fee collections associated with

[3]Unlike centralized computing, "edge computing" pertains to a technology paradigm where vast quantities of data are processed locally by distributed devices, instead of being relayed to a centralized cloud for computation. This local processing approach promises expedited data handling, diminished latency, and improved system efficiency.

legal parking. At that time, Company W set up a system similar to the one depicted in Figure A1, though its function was focused on fee collection. The road cameras captured the moments a vehicle entered and exited a parking position. This data would then be utilized by the traffic police to determine the parking fee. Subsequently, an electronic bill would be sent to the vehicle owner via a mobile app and could be paid with electronic payment. As a result, when introducing AI-automated enforcement in July 2019, Company W did not need to install new roadside cameras in District X. Instead, they merely upgraded the existing devices to incorporate the new function.

This setup enables us to pursue two sets of supplementary analyses, as illustrated in Figure A3. First, we can use the data on legal parking to study the ripple effects on lawful behavior, in addition to the deterrence effects on illegal behavior. Theoretically, enhanced detection and deterrence of illegal actions should promote legal behavior. For severe crimes such as assault and robbery, improved policing should not only suppress the crimes but also make citizens feel safer and more inclined to engage in outdoor activities. In our context, AI-automated enforcement could deter illegal parking, leaving more time and space available for legal parking. We can empirically test this using the records of legal parking.

Second, remember that AI-automated enforcement of illegal parking was rolled out in two phases. In the first phase, while the devices in District X had been upgraded to detect illegal parking, the traffic police did not act upon the photographic evidence. This allows us to perform a falsification test with the device upgrade (before the official adoption) as a pseudo-intervention. If drivers had noticed the camera upgrade and suspected tighter enforcement at this stage, we should have observed some ripple effects on legal parking even before the police adopted the technology for regulating illegal parking. Note that we can only perform such a falsification test on legal parking behavior, because we need to observe the behavior *before* the pseudo-intervention. This works for legal parking, as Company W was involved in parking fee collection before upgrading the devices for enforcement against illegal parking. It does not work for illegal parking because we cannot observe the extent of illegal parking until Company W has upgraded the devices for enforcement.

Unlike District X, which introduced AI-automated collection of parking fees in 2018 and started to automatically regulate illegal parking in 2019, District Y procured both services from the outset in February 2020. As a result, Company W installed two distinct sets of cameras: one for recording legal parking and another for detecting illegal parking. Consequently, for District Y, data on legal and illegal parking cannot be integrated using device IDs. Thus, when we examine the ripple effects of AI-automated enforcement on parking facility efficiency in Section 5.1, our sample is limited to District X. The analysis of illegal parking behavior uses data from both districts.

Lastly, as summarized in Section 1, the pricing of parking has received great attention in extant studies (e.g., Krishnamurthy & Ngo, 2020). However, parking prices in City Z have remained unchanged since 2018, which rules out explicit price changes as a potential explanation of the subsequent empirical findings in both legal and illegal parking.

3 Data and Identification Strategy

3.1 Data and Sample Construction

In this paper, we have compiled data from three sources.

Records of Illegal Parking. – Company W provides us access to all records of illegal parking in Districts X and Y. Each record includes: (i) the precise time when the illegal parking was detected, accurate to the minute; (ii) the parking lot where the violation occurred; (iii) the license plate number of the illegally parked car;[4] (iv) the vehicle brand and car model, inferred from the images captured by Company W; and (v) whether that car had legally parked in the same parking lot within the past 7 (or 14) days.

Two data features are worth clarifying. First, the primary unit of monitoring is a parking lot (referred to as *Dian Wei* in Chinese). This is defined as the area within the purview of a single monitoring camera. Consequently, in our subsequent analysis, the original records will be aggregated at this parking lot level. Second, car brand and model are deduced using an AI algorithm that Company W developed in image recognition and a third-party database of real car shapes by brand and model. Like any AI tool, this algorithm is not infallible. Based on the internal evaluation of Company W, the algorithm has an accuracy rate exceeding 95%.

Parking Lot Information. – Company W maintains cross-sectional data on the parking lots monitored by its cameras. Specifically, for each parking lot, we can access its geographical coordinates (i.e., longitude and latitude) and the calendar date when the traffic police decided to use that lot's monitoring images as evidence of illegal parking. This date is termed by the traffic police as the "activation date" for that specific lot. Additionally, we utilize the geographical coordinates to determine the distance from each parking lot to its closest transportation station (bus or subway) and shopping malls. These distances, indicative of local parking demand and traffic flow, will be referenced as "parking lot attributes" throughout this paper.

Efficiency of Parking Facilities. – As mentioned in Section 2.2, before integrating the capability to detect illegal parking, Company W first employed an AI algorithm to collect information on legal parking in District X, which was then used by the traffic police to calculate parking fees. With this data available, Company W can determine the number of unique parking episodes and the total hours occupied for each parking position every day.[5] For illustration, consider that *Car A* parked in a position for 1 hour, followed by *Car B* occupying the same position for another 2 hours, and then *Car A* returned to use the same position for an additional 1.5 hours. If no other parking occurred in that position on that day, the total unique parking episodes would be 3, and the position had been occupied for 4.5 hours. We then calculate the average of these two variables at the parking-lot-day level, as the details of illegal parking records are at this granularity.

As noted before, we can only merge data on legal and illegal parking for District X since those two datasets are collected by the same sets of monitoring cameras in District X only. For

[4]For privacy reasons, plate numbers have been scrambled, and we do not have access to the exact numbers.

[5]Throughout this paper, the term "parking position" denotes space for ONE car, and each parking lot may consist of multiple parking positions.

all the following discussions regarding the utilization efficiency of parking lots, the empirical analysis would be restricted to the sample of District X.

To summarize, we construct two separate datasets for empirical analyses. First, for illegal parking, we aggregate illegal parking records at the parking-lot-day level and merge them with cross-sectional lot attributes. It is composed of two sub-panels: one covers 521 parking lots in District X (493 treated and 28 never-treated), spanning from July 1, 2019 to February 28, 2021; the other covers 447 parking lots in District Y (259 treated and 188 never-treated), spanning from February 1, 2020 to February 28, 2021.

Second, for legal parking, we combine data on parking facility efficiency and lot attributes of 1,098 parking lots in District X. The resulting parking-lot-day panel spans from January 1, 2019 to February 28, 2021. Please note that this data has more than doubled the number of parking lots as compared to the illegal parking data of District X because a number of parking lots have not been upgraded to detect illegal parking during our sample period.

Table A1 details how Districts X and Y have a staggered lot-by-lot adoption of AI-automated enforcement against illegal parking, highlighting their differences in data mergeability. Figure A4 graphically summarizes the progressive expansion of AI-automated enforcement in both districts.

3.2 Identification Strategy

Our econometric strategies are deeply rooted in the background and data as described in Section 2. In particular, the traffic police adopted the image evidence for issuing tickets in a staggered way, which provides a perfect context for a DID model. Additionally, illegal parking records are accurate to the minute and police activation time is accurate to specific calendar dates. This allows us to use Regression Discontinuity Design (RDD) where the number of days relative to the time of treatment serves as the running variable.[6]

Following Li et al. (2020), to evaluate how AI-automated enforcement affects illegal parking, we have employed both DID and RDD to strengthen the credibility of the empirical analysis and make full use of the data. To save space, we focus on the DID model in the main text, and relegate the details of RDD to Online Appendix C. With regard to the DID model, we estimate the following equation:

$$Y_{it} = \beta PostTreat_{it} + \delta^{'} X_{it} + \lambda_i + \theta_t + \varepsilon_{it}\,, \tag{1}$$

where Y_{it} denotes the outcome variable such as the logarithmic number of illegal parking at parking lot i on date t. $PostTreat_{it}$ is a dummy variable that takes the value of 1 if records from parking lot i have already been adopted for enforcement by the traffic police on date t and 0 otherwise. Its coefficient β is the key coefficient that captures the causal impact of AI-automated enforcement on illegal parking behavior. In the remainder of this article, "treatment group" refers to the parking lots where the traffic police have ever adopted AI-automated en-

[6]For more about RDD using relative time as the running variable, see Hausman and Rapson (2018) for a comprehensive review and Auffhammer and Kellogg (2011), Gallego et al. (2013), and Anderson (2014) for application examples.

forcement during the sample period. By contrast, "control group" refers to the parking lots where monitoring cameras have also been upgraded to detect illegal parking but the related images have never been adopted for enforcement by the traffic police. Our DID model essentially compares the treatment and control groups, after controlling for parking lot fixed effects (λ_i) and date fixed effects (θ_t), and accounting for different activation dates for different treated lots.

To address the possibility that different lots may face different time-varying demands for parking, we use X_{it} to denote two sets of time-varying controls. The first is lot-specific polynomial trends. We first compute a linear trend that takes the value of 1, 2, 3, ... since July 1, 2019, for lots in District X (correspondingly, the starting date is February 1, 2020, for lots in District Y). The linear trend, along with its quadratic and cubic terms, is further interacted with the lot fixed effects. Second, we compute the distance from each parking lot to its nearest transportation stations (bus or subway) and shopping malls. We interact these two parking lot attributes with date fixed effects to allow for more flexible daily shocks experienced by different parking lots. ε_{it} is the error term.

Even with these additional controls, the causal interpretation of β relies on the assumption of parallel pre-treatment trend. In the interview with the traffic police, they claimed that their decisions on activation dates were random without any explicit rules. To further test the assumption of pre-trend and display the evolution of treatment effects, we estimate the following event study model:

$$Y_{it} = \sum_{k=-12}^{12} \gamma_k A_{ik} + \delta^{'} X_{it} + \lambda_i + \theta_t + \varepsilon_{it} \,. \tag{2}$$

To increase the statistical power of point estimates and examine data patterns in a period as long as possible, we collapse every two weeks (i.e., 14 days) as one bin. A_{ik} is a series of dummies indicating bins -12, -11, ..., -1 before the treatment for parking lot i, and bins 1, 2, ..., 12 after the treatment. For example, $A_{i,-3}$ denotes 5–6 weeks before the police adopted AI-automated enforcement on parking lot i. Estimates of $\{\gamma_k\}_{k<0}$ indicate whether the outcome variable of treated parking lots varies similarly to that of the control group in a period as early as 6 months before the treatment; estimates of $\{\gamma_k\}_{k\geq 0}$ represent the distribution of the treatment effects for each bin after the treatment.

4 Empirical Results

4.1 Baseline Results

Table B1 reports the summary statistics of illegal and legal parking in our analysis samples. The unit of observation is parking-lot-day. A parking lot is treated if the local traffic police has ever adopted AI-automated enforcement for illegal parking in our sample period. Otherwise, the parking lot is in the control group. Within treated lots, we observe a sharp decline in the number of illegal parking episodes, from an average of 1.876 per lot-day before treatment to 0.671 after treatment. Consistently, the likelihood of having any illegal parking in a lot-day has dropped from 38.8% before treatment to 25.6% after treatment. Admittedly, the treated

parking lots were more likely to observe illegal parking than control lots before the treatment, which highlights the importance of controlling for parking lot fixed effects, calendar date fixed effects, and the time-varying impacts of parking lot attributes before we attribute any before-after difference as the causal effect of AI-automated enforcement.

As for legal parking behavior, our sample is limited to District X but has more lot-day observations than the sample of illegal parking because many parking lots in District X have never upgraded the devices for illegal parking enforcement but they have used Company W's services for parking fee collection. Table B1 shows that District X faces high demand for legal parking: each parking position serves more than 5 parking episodes per day and is occupied for more than 10 hours per day. Within treated parking lots, both statistics increase substantially after the treatment: from 5.512 to 6.206 for the count of legal parking episodes and from 9.275 to 11.769 for occupied hours. Compared to control lots, the treated lots had more episodes but fewer occupied hours before the treatment. It is unclear what types of selection may drive such differences, but our regression analyses will use extensive controls to address the potential selection.

Table 1 reports the baseline results from estimating equation (1). To show its sensitivity to model selection, we add the controls and time trends in a gradual way. We begin with the simplest DID setting in Column (1) of Panel A, where we only control for parking lot and calendar date fixed effects. The estimated coefficient implies that the number of illegal parking within each parking lot per day will decrease by 13.4% ($= e^{0.126} - 1$) on average after the AI-automated detection was adopted by traffic police.

In Column (2), we further include the interactions between parking lot fixed effects and linear, quadratic, and cubic terms of a time trend. The trend takes the value of 1, 2, 3, ... since July 1, 2019, for parking lots in District X (correspondingly, since February 1, 2020, for parking lots in District Y). The inclusion of those terms allows different parking lots to exhibit intrinsic trends in the occurrence of illegal parking. In Column (3), we further add the interactions between date fixed effects and parking lot attributes, which are the distances between each parking lot and its nearest transportation stations and shopping malls. Those two distance measures should help to capture the differences in the traffic flows surrounding parking lots because commuting demand is usually stronger for locations near hubs like stations and shopping malls. Including their interactions with date fixed effects can account for the heterogeneous shocks experienced by parking lots facing different traffic flows.

As shown in Panel A of Table 1, the magnitude of our key coefficient β decreases by about a half (from 0.126 to about 0.059) as we gradually control for more complex time trends and lot attributes. However, even for the most stringent specification in Column (3) (referred to as the "benchmark" specification), we still find illegal parking decreases by 6% ($= e^{0.059} - 1$) after the adoption of AI-automated enforcement, which is both economically and statistically significant.

In Panel B of Table 1, the outcome variable is changed into a dummy indicating whether there is any illegal parking detected at parking lot i on date t. Again, the findings are robust regardless of how many time-varying controls we have on the right-hand side. In the most

stringent specification (Column 3), AI-automated enforcement leads to a 3.3 percentage point drop in the occurrence of any illegal parking in a parking-lot-day, which is a significant improvement from the pre-treatment mean of 38.8 percentage points.

Before delving deeper into the DID model, it is worth noting that the RDD estimation produced remarkably similar empirical findings, suggesting that AI-automated enforcement effectively deters illegal parking. While there may be differences in coefficient magnitudes, arriving at consistent findings through another parallel identification strategy strengthens the credibility of the DID model. Estimates, along with robustness checks for the RDD model, can be found in Online Appendix C.

4.2 Event Study

As outlined in Section 3.2, we estimate equation (2) to conduct an event study, which tests the assumption of parallel pre-trend within 12 bins before the treatment and displays the distribution of treatment effects for 12 bins post the treatment. Recall that we have collapsed every two weeks as one bin, so the test allows us to examine the pre-trend during a time period as long as 6 months. For the identification of the model, the period just before the police's adoption ($k = -1$) is dropped as the omitted category.

Figure 1 presents the results of the event study. The graphic evidence indicates that $\{\gamma_k\}_{k<0}$ are all statistically insignificant. This implies that, as early as 6 months before the adoption of AI enforcement by traffic police, parking lots in treatment and control groups have shown no obvious discrepancy in observed illegal parking. In addition, we find that the AI-automated enforcement has immediate and persistent deterrence impact on illegal parking, as reflected by the significant estimates of $\{\gamma_k\}_{k\geq 0}$.

4.3 Robustness Checks

We further conduct three robustness checks for the DID model. The corresponding results are reported in Table B2.

First, while the pre-trend test suggests that treated and control parking lots follow similar trends in illegal parking (after all the controls in equation (1)), concerns may still arise about the comparability of parking lots between these groups. Adhering to the method outlined by Li et al. (2016), we exclude all parking lots from the control group, focusing our analysis solely on a before/after comparison within the treated parking lots. Because their treatment activation dates were staggered, we can still identify the DID model. Column (1) of Table B2 indicates only a slight shift in the point estimate of β (from -0.059 to -0.056), even though the sample size is reduced significantly. This strengthens our confidence that the baseline findings are unlikely driven by inherent differences between the two groups of parking lots.

Second, as shown by the latest literature on DID and event studies, OLS estimation of the DID model may suffer from potential bias due to heterogeneous treatment effects or the "negative weights" problem (Goodman-Bacon, 2021). As a response, in addition to the classic two-way fixed effects (TWFE) OLS estimation, we have estimated equation (1) using methods proposed by Borusyak et al. (2021) and Cengiz et al. (2019). The corresponding results are reported in Columns (2) and (3) of Table B2, respectively. Results suggest that the estimated de-

terrence effect on illegal parking remains stable in magnitude and is statistically significant at at least the 10% level. Coefficients obtained from alternative estimators indicate that potential bias of classical TWFE OLS should not be a key concern for our DID model.[7]

Third, we conduct analysis using a Poisson regression, which is commonly employed when the dependent variable represents count data. Specifically, we use the model developed by Correia et al. (2020) and Correia et al. (2019), which allows for intricate fixed effects and interaction terms within the Poisson framework. As illustrated in Column (4) of Table B2, the Poisson regression yields findings consistent with those from linear estimation, confirming that AI-automated enforcement has effectively curtailed illegal parking.[8] The estimated coefficient suggests that the count of illegal parking during the post-treatment period is approximately 79% ($\approx e^{-0.235}$) of its pre-treatment level. This translates to a 21% reduction in illegal parking, a figure that stands out even more when juxtaposed with results from our benchmark model.

5 Further Analysis

5.1 Ripple Effects on Parking Efficiency

As introduced in Section 2.2, legal parking data in District X help us to go beyond the deterrence impact on illegal parking and delve into the ripple effects on parking efficiency. A serious look into legal behavior is highly relevant because the ultimate goal of parking rules (along with AI-automated enforcement) is to ensure that cars are parked in an organized way, which would further contribute to the order and smoothness of traffic (Inci, 2015).

To empirically study this, we combine data on parking facility efficiency and AI-automated enforcement in District X. The number of unique parking episodes and occupied hours per parking position per day were averaged at the parking-lot-day level and then used as outcome variables in equation (1). As shown in the first two columns of Table 2, AI-automated enforcement of illegal parking has effectively increased the number of unique parking episodes by 0.267 on a position-day basis. Though it only implies an increase of about 5% of the sample mean, it is still economically significant considering that there are more than 70,000 curbside parking positions in City Z (as of 2017). On the contrary, we do not observe a similar improvement in occupied hours, as shown in Column (2).

The results in Columns (1) and (2) make economic sense. On the one hand, because illegal parking behavior (such as double parking) has been effectively deterred by AI enforcement, more space is available for legal parking at more times during the day, which could improve the turnover rate on each parking position. That is why overall parking episodes increase significantly. At the same time, the length of time that a driver needs roadside parking on a specific street is likely exogenous, hence the drop in illegal parking does not necessarily induce

[7] There are many alternative estimators such as De Chaisemartin and d'Haultfoeuille (2020) or Sun and Abraham, 2021, which are systematically reviewed in De Chaisemartin and d'Haultfoeuille (2022). However, not all new methods apply to this paper, which is conducted with an unbalanced panel and highly sparse treatment times. Even for the two methods adopted here, we have to sacrifice some dimensions of the benchmark model (e.g., dropping the linear, quadratic, and cubic trends specific to parking lots) to guarantee the convergence of the estimation.

[8] As in the second robustness check, we have to drop the linear, quadratic and cubic trends specific to parking lots in order to achieve convergence in the Poisson regression.

the driver to prolong the parking time or initiate unnecessary parking (like roadside parking late at night). This explains why the total occupied hours of parking lots have changed little.

Lastly, recall that the AI-automated enforcement was implemented in two phases. During the first phase, the traffic police did not utilize the picture evidence detected by AI. To perform a falsification test, we limit our sample to pre-adoption observations and create a false treatment variable representing the upgrading of Company W's devices in District X. Results from Columns (3) and (4) of Table 2 show no ripple effects from AI-automated enforcement when it is not in active use. This suggests that the gains in parking efficiency do not exist until AI-automated enforcement is activated, which reinforces the linkage between the enforcement of illegal behavior and the facilitation of legal behavior.

5.2 Distributional Effects

As emphasized by Acemoglu (2021), when considering the effects of AI, it should not be perceived merely as a straightforward common shock. Instead, the distribution of effects across different groups of people – referred to as distributional effects – is of paramount importance. In this subsection, we delve into the potential varied impacts of AI-automated enforcement across several categories of drivers.

First, we investigate whether people of different economic status may respond differently to AI-automated enforcement. We do not have any information on driver income or wealth, but thanks to the powerful AI image recognition algorithm developed by Company W, we are able to construct a proxy that is informative about the driver's economic status. Specifically, for each piece of illegal parking record, Company W has utilized its AI algorithm to perform image recognition and infer the specific car brand and model according to the car's appearance (e.g., Audi A6, BMW 3).[9] Based on this information, we further collect the price range (i.e., the maximum and minimum recommended retail prices) for each car model from the website AutoHome, which specializes in collecting price quotes for different cars.[10] For each car model, we first calculate its mean recommended retail price and then categorize it as either *Expensive* or *Cheap* according to whether its mean price is above the median of all car models. In this way, we decompose parking-lot-day illegal parking into two subgroups and re-compute the total numbers as outcome variables, respectively. Illegal parking records that the AI algorithm fails to identify are categorized into a residual group labeled as *Unidentified*.

The first three columns of Table 3 report corresponding results. Columns (1) and (2) show that drivers of cheap cars are much more sensitive to AI-automated deterrence on illegal parking, which echoes extant opinions that AI may affect people of heterogeneous economic statuses differently.[11] If we interpret the enhanced law enforcement automated by AI as an in-

[9] The training data set for the algorithm comes from another company which specializes in collecting and posting car-related information. Company W combined this third-party database with its AI algorithm to identify car brands and models, so it is immune from the potential bias of algorithm learning (e.g., someone may be concerned that the algorithm may be more and more effective in identifying cheap brands because of a larger data volume).

[10] The website is https://www.autohome.com.cn/.

[11] A bootstrap test of the difference between coefficients in Columns (1) and (2) shows that the difference is statistically significant at a 1% significance level. Panel A of Figure B1 plots the distribution of coefficients obtained in the bootstrap test. Similarly, we have also conducted bootstrap tests for the other three dimensions of distributional effects, and corresponding results are plotted in Panels B – D of Figure B1.

crease in the "cost" of illegal parking, it is conceivable that people of higher economic status are less responsive to the increased cost. In the first four columns of Table B3, we use the minimum or maximum (instead of mean) recommended retail prices to classify car models as expensive or cheap and find that results are almost the same as Table 3.

Second, we examine whether drivers' responses are different during peak hours, when they may face more binding constraints on the time and location of parking. For example, during peak hours, the driver may need to park the car and hurry to work, which means that he or she may have no choice but to park illegally. If that is the case, drivers should be less responsive to AI-automated regulation of illegal parking during peak hours.

To empirically test this conjecture, we categorize parking-lot-day illegal parking into two groups based on whether it occurs during peak or non-peak hours. Peak hours are defined as 7:00 – 9:00 and 17:00 – 20:00, in line with the patterns of tidal traffic flow in City Z. As shown in Table 3, the magnitude of coefficient in Column (4) is significantly smaller than Column (5). It is consistent with our expectation that when a driver has a strong preference for the time and location of parking, the deterrence effect of AI-automated law enforcement is largely weakened.

Finally, we investigate how local affinity affects people's responses when they face more effective enforcement against misbehavior. According to the first Chinese character on licensing plate numbers, we can identify whether the car is registered in City Z or not. The first two columns of Table 4 show that local drivers, who have to adapt themselves to the new policy, respond more intensively to AI-automated enforcement. Non-local drivers may demonstrate weaker responses for either intentional reasons (for instance, they do not care about occasional illegal parking in another city) or unintentional ignorance (for instance, they do not know the AI-automated enforcement in City Z). We cannot distinguish those two potential mechanisms due to data limitations.

Furthermore, we categorize local drivers into regular and occasional commuters based on whether they park at the parking lot regularly. For each parking lot, some drivers may visit it regularly (e.g., driving to work every day) while others may rarely drive there. Will regular and occasional commuters behave differently? To examine this question, we define *Regular* or *Occasional* commuters according to whether the car plate number has emerged in legal parking records of the parking lot in the past 7 days, and then use the newly constructed measures as the outcome variable. Empirical results in Columns (3) and (4) of Table 4 indicate that occasional commuters exhibit a stronger response to AI-automated enforcement. One potential explanation is that regular commuters are already familiar with nearby parking lots and therefore have limited room to improve compliance with parking rules. Another possibility is that the focal location provides substantial convenience to regular commuters, which makes them less likely to switch somewhere else even if misbehavior at this location faces harsher enforcement.[12]

[12]In Table B3, we try to extend the time window for defining regular and occasional commuters from 7 to 14 days. As shown in Columns (5) and (6), though the coefficient difference shrinks a little, those two kinds of commuters still exhibit significantly heterogeneous responses to AI-automated enforcement.

6 Conclusion

In recent years, the rapid proliferation of AI has been unmistakable. While AI's contributions to production activities in private sectors are well-acknowledged, its role in public administration has largely been overlooked. To fill this gap, we utilize detailed data from a private Chinese company to empirically assess the impact of AI-automated law enforcement on illegal and legal parking. Our analyses show that AI-automated enforcement significantly curtails illegal parking and improves the efficiency of parking facilities. We also find significant distributional effects by drivers' economic status, driving time, local affinity, and commute regularity.

While our work highlights the benefits of AI-automated enforcement, data limitation prevents us from addressing many potential costs of it. For example, by definition, the AI technology employed in parking detection focuses on the vehicle rather than the driver. It is imaginable that other AI technologies could be more driver-focused – for example, AI may detect whether the driver has worn the seat belt while driving or whether the car has been occupied by too many passengers – and therefore be more intrusive to the driver or passenger privacy. Similarly, the parking lots in our data are all roadside, which arguably are more public and less privacy-sensitive than garages. How to trade off privacy concerns against the benefits of AI-automated enforcement is an important question beyond our paper. The answer may differ by the situation subject to law enforcement and the capability of the AI technology.

In addition, the technology studied in this paper treats all vehicles in the same way, which minimizes human error and human bias. However, it also leaves less room for pardoning minor violations without a fine or allowing extraordinary situations. In theory, these exceptions can be addressed by appeal, but that may require the ticket receiver to spend more time and effort than explaining it to a human patrol on the spot. For this reason, it remains an open question whether all citizens would prefer the new precise-but-cold system to the old system with more human touches. How to deal with errors and exceptions in the AI-powered system and how to make the system more humane deserve future research.

References

Acemoglu, D. (2021). Harms of AI. *National Bureau of Economic Research Working Paper*.

Anderson, M. L. (2014). Subways, Strikes, and Slowdowns: The Impacts of Public Transit on Traffic Congestion. *American Economic Review*, *104*(9), 2763–2796.

Arnott, R. (2011). Parking Economics. In A. De Palma, R. Lindsey, E. Quinet, & R. Vickerman (Eds.), *A Handbook of Transport Economics* (pp. 726–743). Edward Elgar Publishing.

Arnott, R., & Inci, E. (2006). An Integrated Model of Downtown Parking and Traffic Congestion. *Journal of Urban Economics*, *60*(3), 418–442.

Auffhammer, M., & Kellogg, R. (2011). Clearing the Air? The Effects of Gasoline Content Regulation on Air Quality. *American Economic Review*, *101*(6), 2687–2722.

Bauernschuster, S., & Rekers, R. (2022). Speed Limit Enforcement and Road Safety. *Journal of Public Economics*, *210*, 104663.

Borusyak, K., Jaravel, X., & Spiess, J. (2021). Revisiting Event Study Designs: Robust and Efficient Estimation. *arXiv preprint arXiv:2108.12419*.

Calonico, S., Cattaneo, M. D., & Titiunik, R. (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica*, *82*(6), 2295–2326.

Cattaneo, M. D., Idrobo, N., & Titiunik, R. (2019). *A Practical Introduction to Regression Discontinuity Designs: Foundations*. Cambridge University Press.

Cengiz, D., Dube, A., Lindner, A., & Zipperer, B. (2019). The Effect of Minimum Wages on Low-Wage Jobs. *Quarterly Journal of Economics*, *134*(3), 1405–1454.

Chalfin, A., & McCrary, J. (2017). Criminal Deterrence: A Review of the Literature. *Journal of Economic Literature*, *55*(1), 5–48.

Correia, S., Guimarães, P., & Zylkin, T. (2019). Verifying the Existence of Maximum Likelihood Estimates for Generalized Linear Models. *arXiv preprint arXiv:1903.01633*.

Correia, S., Guimarães, P., & Zylkin, T. (2020). Fast Poisson Estimation with High-Dimensional Fixed Effects. *Stata Journal*, *20*(1), 95–115.

De Chaisemartin, C., & d'Haultfoeuille, X. (2020). Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review*, *110*(9), 2964–2996.

De Chaisemartin, C., & d'Haultfoeuille, X. (2022). Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey. *National Bureau of Economic Research Working Paper*.

Dušek, L., & Traxler, C. (2022). Learning from Law Enforcement. *Journal of the European Economic Association*, *20*(2), 739–777.

Gallego, F., Montero, J.-P., & Salas, C. (2013). The Effect of Transport Policies on Car Use: Evidence from Latin American Cities. *Journal of Public Economics*, *107*, 47–62.

Goodman-Bacon, A. (2021). Difference-in-Differences with Variation in Treatment Timing. *Journal of Econometrics*, *225*(2), 254–277.

Hasker, K., & Inci, E. (2014). Free Parking for All in Shopping Malls. *International Economic Review*, *55*(4), 1281–1304.

Hausman, C., & Rapson, D. S. (2018). Regression Discontinuity in Time: Considerations for Empirical Applications. *Annual Review of Resource Economics*, *10*, 533–552.

Hernández, P., & Neeman, Z. (2022). How Bayesian Persuasion Can Help Reduce Illegal Parking and Other Socially Undesirable Behavior. *American Economic Journal: Microeconomics*, *14*(1), 186–215.

Imbens, G., & Kalyanaraman, K. (2012). Optimal Bandwidth Choice for the Regression Discontinuity Estimator. *Review of Economic Studies*, *79*(3), 933–959.

Inci, E. (2015). A Review of the Economics of Parking. *Economics of Transportation*, *4*(1-2), 50–63.

Jabri, R. (2021). Algorithmic Policing. *Working Paper*.

Kapustin, M., Neumann, T., & Ludwig, J. (2022). Policing and Management. *National Bureau of Economic Research Working Paper*.

Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J., & Mullainathan, S. (2018a). Human Decisions and Machine Predictions. *Quarterly Journal of Economics*, *133*(1), 237–293.

Kleinberg, J., Ludwig, J., Mullainathan, S., & Sunstein, C. R. (2018b). Discrimination in the Age of Algorithms. *Journal of Legal Analysis*, *10*, 113–174.

Kobus, M. B., Gutiérrez-i-Puigarnau, E., Rietveld, P., & Van Ommeren, J. N. (2013). The On-Street Parking Premium and Car Drivers' Choice between Street and Garage Parking. *Regional Science and Urban Economics*, *43*(2), 395–403.

Krishnamurthy, C. K. B., & Ngo, N. S. (2020). The Effects of Smart-Parking on Transit and Traffic: Evidence from SFpark. *Journal of Environmental Economics and Management*, *99*, 102273.

Lee, D. S., & Lemieux, T. (2010). Regression Discontinuity Designs in Economics. *Journal of Economic Literature*, *48*(2), 281–355.

Li, P., Lu, Y., & Wang, J. (2016). Does Flattening Government Improve Economic Performance? Evidence from China. *Journal of Development Economics*, *123*, 18–37.

Li, P., Lu, Y., & Wang, J. (2020). The Effects of Fuel Standards on Air Pollution: Evidence from China. *Journal of Development Economics*, *146*, 102488.

Lin, H., & Wang, I. Y. (2015). Competition and Price Discrimination: Evidence from the Parking Garage Industry. *Journal of Industrial Economics*, *63*(3), 522–548.

Ludwig, J., & Mullainathan, S. (2021). Fragile Algorithms and Fallible Decision-Makers: Lessons from the Justice System. *Journal of Economic Perspectives*, *35*(4), 71–96.

Mastrobuoni, G. (2020). Crime is Terribly Revealing: Information Technology and Police Productivity. *Review of Economic Studies*, *87*(6), 2727–2753.

Nissan, E. (2017). Digital Technologies and Artificial Intelligence's Present and Foreseeable Impact on Lawyering, Judging, Policing and Law Enforcement. *AI & Society*, *32*, 441–464.

Polinsky, A. M., & Shavell, S. (2000). The Economic Theory of Public Enforcement of Law. *Journal of Economic Literature*, *38*(1), 45–76.

Shoup, D. C. (2006). Cruising for Parking. *Transport Policy*, *13*(6), 479–486.

Sun, L., & Abraham, S. (2021). Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects. *Journal of Econometrics*, *225*(2), 175–199.

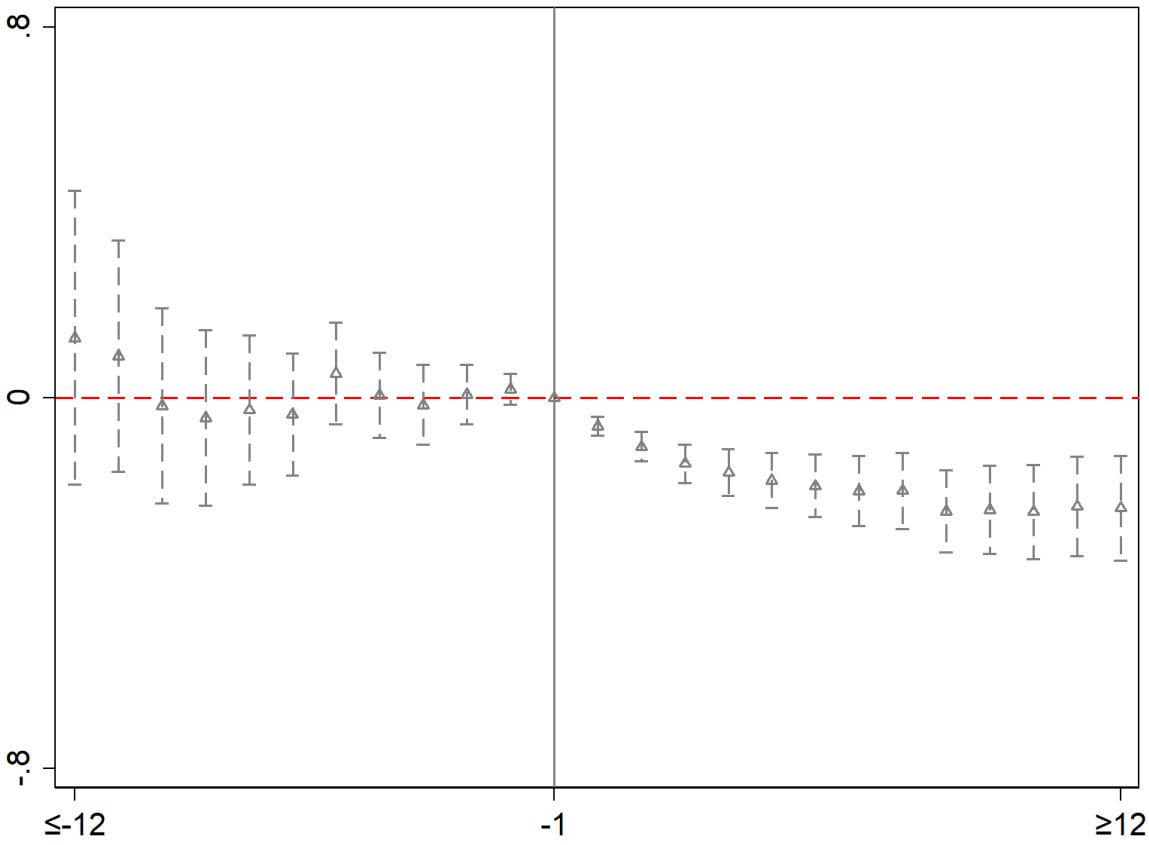
Traxler, C., Westermaier, F. G., & Wohlschlegel, A. (2018). Bunching on the Autobahn? Speeding Responses to a 'Notched' Penalty Scheme. *Journal of Public Economics*, *157*, 78–94.

Venkatraman, V., Richard, C. M., Magee, K., & Johnson, K. (2021). *Countermeasures That Work: A Highway Safety Countermeasure Guide for State Highway Safety Offices, 2020* (tech. rep.). National Highway Traffic Safety Administration.

Zakharenko, R. (2019). The Economics of Parking Occupancy Sensors. *Economics of Transportation*, *17*, 14–23.

Figures and Tables

Figure 1. *Estimates from the Event Study*



Notes: This figure plots point estimates along with 95% confidence intervals from the event study specified in equation (2). We have collapsed every two weeks (14 days) as one bin to increase the statistical power and the period just before the police's adoption ($k = -1$) is dropped as the omitted category.

Table 1. *Impact of AI-Automated Enforcement Against Illegal Parking*

	(1)	(2)	(3)
Panel A	*Log (Illegal Parking)*		
PostTreat	-0.126***	-0.063***	-0.059***
	(0.029)	(0.017)	(0.018)
Adjusted R^2	0.546	0.646	0.647
Panel B	*Dummy for Occurrence of Illegal Parking*		
PostTreat	-0.058***	-0.032***	-0.033***
	(0.015)	(0.011)	(0.011)
Adjusted R^2	0.396	0.463	0.463
Sample	Full	Full	Full
Parking Lot Fixed Effects	YES	YES	YES
Calendar Date Fixed Effects	YES	YES	YES
Parking Lot FE $\times$ (T, T^2, T^3)		YES	YES
Date FE $\times$ Parking Lot Attributes			YES
# of Clusters	968	968	968
# of Observations	431870	431870	431870

Notes: This table reports the impact of AI-automated enforcement on illegal parking. In Column (1), we only control for parking lot and calendar date fixed effects. We further include the interactions between parking lot fixed effects and linear, quadratic, and cubic terms of a time trend in Column (2), and the interactions between date fixed effects and parking lot attributes in Column (3). In Panels A and B, outcome variables are the logarithmic number of illegal parking and the dummy for occurrence of illegal parking, respectively. Standard errors clustered at the parking lot level are reported in parentheses with ***, **, and * denoting significance at the 1%, 5%, and 10% levels.

Table 2. *Ripple Effects on Parking Efficiency and Falsification Test*

	(1)	(2)	(3)	(4)
Outcome Variable	*Number of Unique Parking Episodes*	*Occupied Hours*	*Number of Unique Parking Episodes*	*Occupied Hours*
PostTreat	0.267***	-0.041		
	(0.095)	(0.172)		
PostUpgrade			-0.067	-0.201
			(0.123)	(0.237)
Sample	District X	District X	District X	District X
Parking Lot Fixed Effects	YES	YES	YES	YES
Calendar Date Fixed Effects	YES	YES	YES	YES
Parking Lot FE $\times$ (T, T^2, T^3)	YES	YES	YES	YES
Date FE $\times$ Parking Lot Attributes	YES	YES	YES	YES
Adjusted R^2	0.528	0.526	0.506	0.512
# of Clusters	1098	1098	1098	1098
# of Observations	847934	847934	618305	618305

Notes: This table reports results from estimation using the data on legal parking in District X. In Columns (1) and (2), we estimate the ripple effects of AI-automated enforcement on parking efficiency. In Columns (3) and (4), we limit the sample to pre-adoption observations and create a false treatment variable representing the upgrading of Company W's devices in District X to conduct the falsification test. Standard errors clustered at the parking lot level are reported in parentheses with ***, **, and * denoting significance at the 1%, 5%, and 10% levels.

Table 3. *Distributional Effects Across Economic Status and Time of Parking*

Outcome Variable	(1)	(2)	(3)	(4)	(5)
	Log (Illegal Parking)				
	Expensive	Cheap	Unidentified	Peak Hours	Non-Peak Hours
PostTreat	-0.031***	-0.053***	-0.011	-0.038***	-0.047***
	(0.010)	(0.012)	(0.007)	(0.010)	(0.014)
Sample	Full	Full	Full	Full	Full
Parking Lot Fixed Effects	YES	YES	YES	YES	YES
Calendar Date Fixed Effects	YES	YES	YES	YES	YES
Parking Lot FE $\times$ (T, T^2, T^3)	YES	YES	YES	YES	YES
Date FE $\times$ Parking Lot Attributes	YES	YES	YES	YES	YES
Adjusted R^2	0.528	0.560	0.433	0.537	0.587
# of Clusters	968	968	968	968	968
# of Observations	431870	431870	431870	431870	431870

Notes: This table reports the distributional effects of AI-automated enforcement through the lens of economic status and parking time. In the first three columns, we categorize cars as either *Expensive* or *Cheap* according to their prices (and *Unidentified* for those not identified by the algorithm). In Columns (4) and (5), we categorize parking-lot-day illegal parking into two groups based on whether it occurs during peak or non-peak hours, where peak hours are defined as 7:00 – 9:00 and 17:00 – 20:00. Standard errors clustered at the parking lot level are reported in parentheses with ***, **, and * denoting significance at the 1%, 5%, and 10% levels.

Table 4. *Distributional Effects Across Local Affinity and Commuting Frequency*

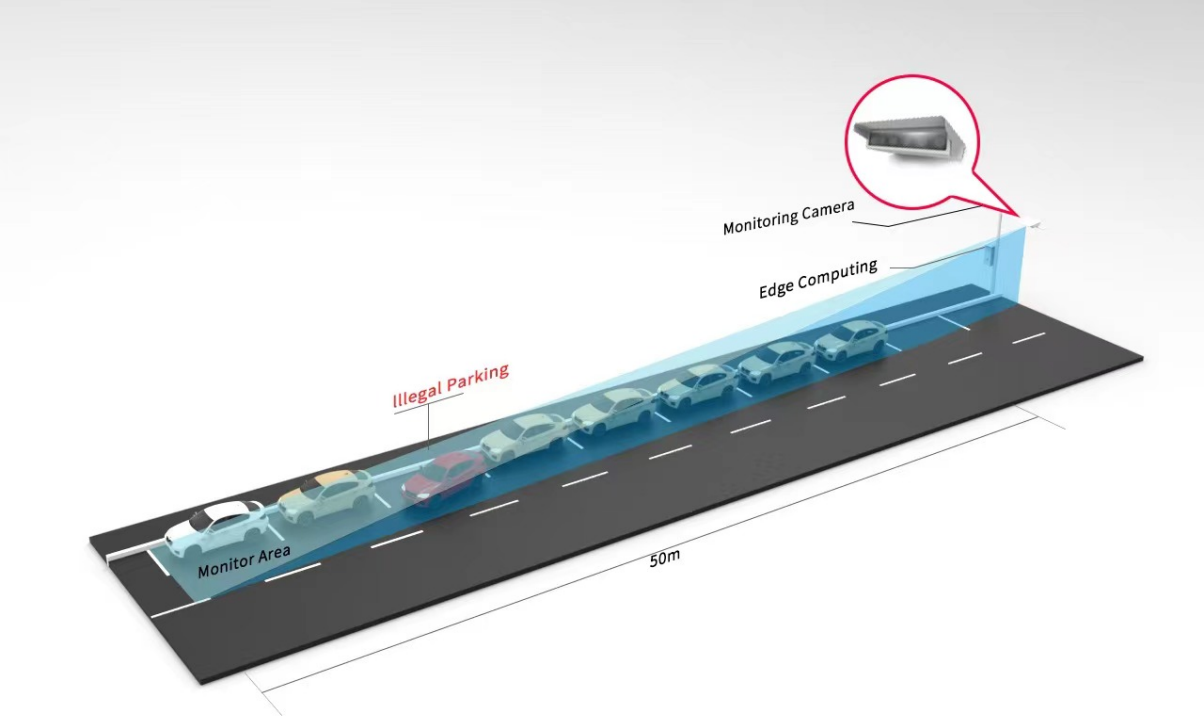
Outcome Variable	(1)	(2)	(3)	(4)
			Log (Illegal Parking)	
	Non-Local	Local	Local & Regular	Local & Occasional
PostTreat	-0.018**	-0.056***	-0.013***	-0.052***
	(0.007)	(0.016)	(0.005)	(0.015)
Sample	Full	Full	Full	Full
Parking Lot Fixed Effects	YES	YES	YES	YES
Calendar Date Fixed Effects	YES	YES	YES	YES
Parking Lot FE $\times$ (T, T^2, T^3)	YES	YES	YES	YES
Date FE $\times$ Parking Lot Attributes	YES	YES	YES	YES
Adjusted R^2	0.504	0.623	0.248	0.619
# of Clusters	968	968	968	968
# of Observations	431870	431870	431870	431870

Notes: This table reports the distributional effects of AI-automated enforcement through the lens of local affinity and commuting frequency of drivers. In Columns (1) and (2), we divide illegal parking records into two groups according to whether the car is registered locally or not. In Columns (3) and (4), we first restrict to cars locally registered and then define *Regular* or *Occasional* commuters according to whether the car plate number has emerged in legal parking records of the parking lot in the past 7 days. Standard errors clustered at the parking lot level are reported in parentheses with ***, **, and * denoting significance at the 1%, 5%, and 10% levels.

Online Appendix

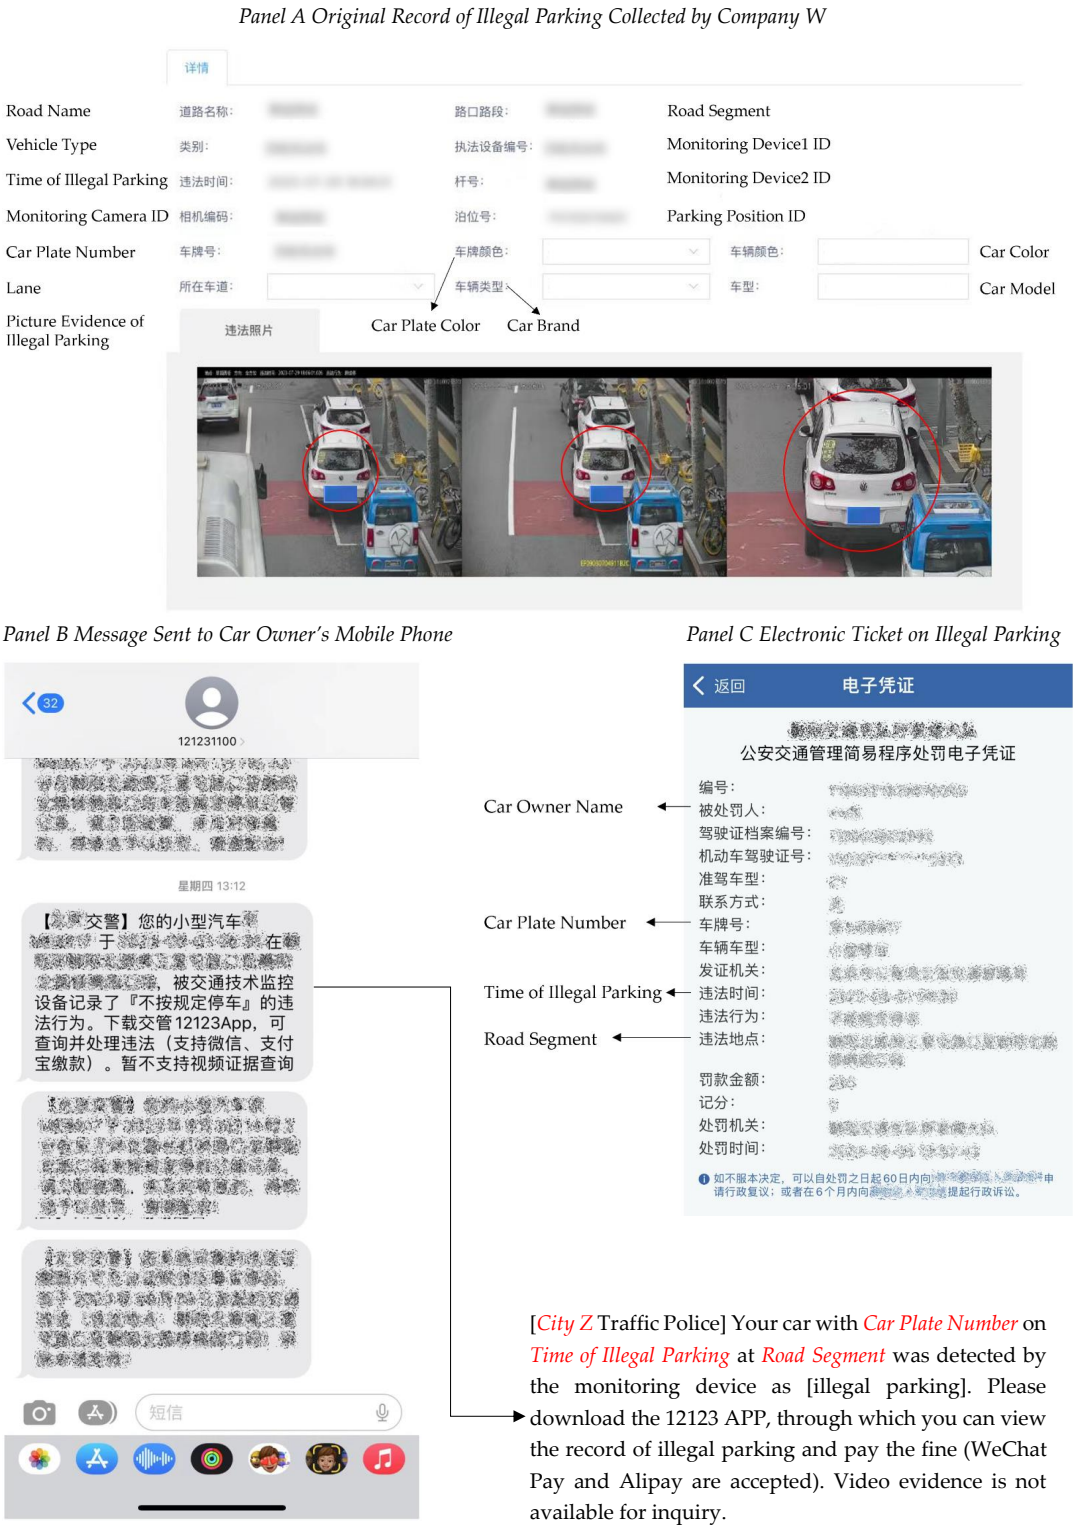
Appendix A Details about AI-Automated Enforcement

Figure A1. *Outlook of the AI-Automated Monitoring System*



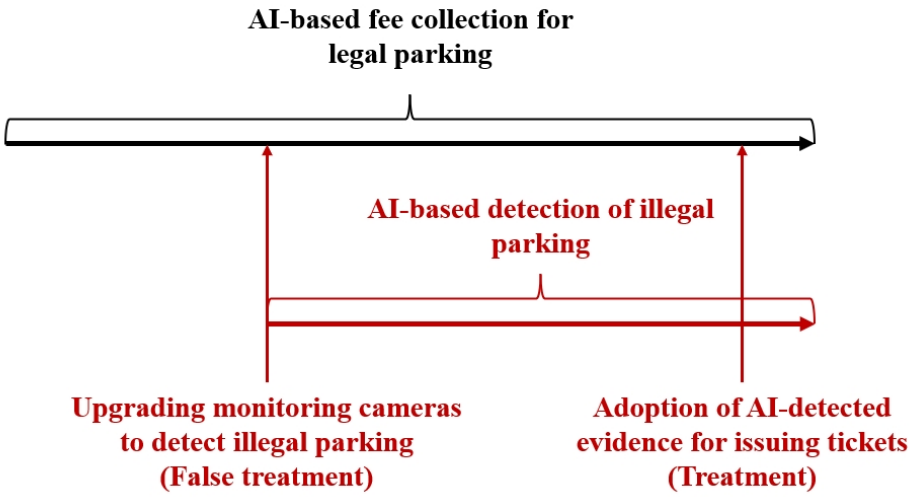
Notes: This figure provides a visual representation of the appearance of the AI-automated monitoring system. The monitoring camera is mounted atop a post and the monitored area depicted in this image encompasses ten parking positions, spanning a total length of 50 meters. The edge computing device is used for conducting AI-powered image recognition locally.

Figure A2. *The Flow Chart of AI-Automated Enforcement*



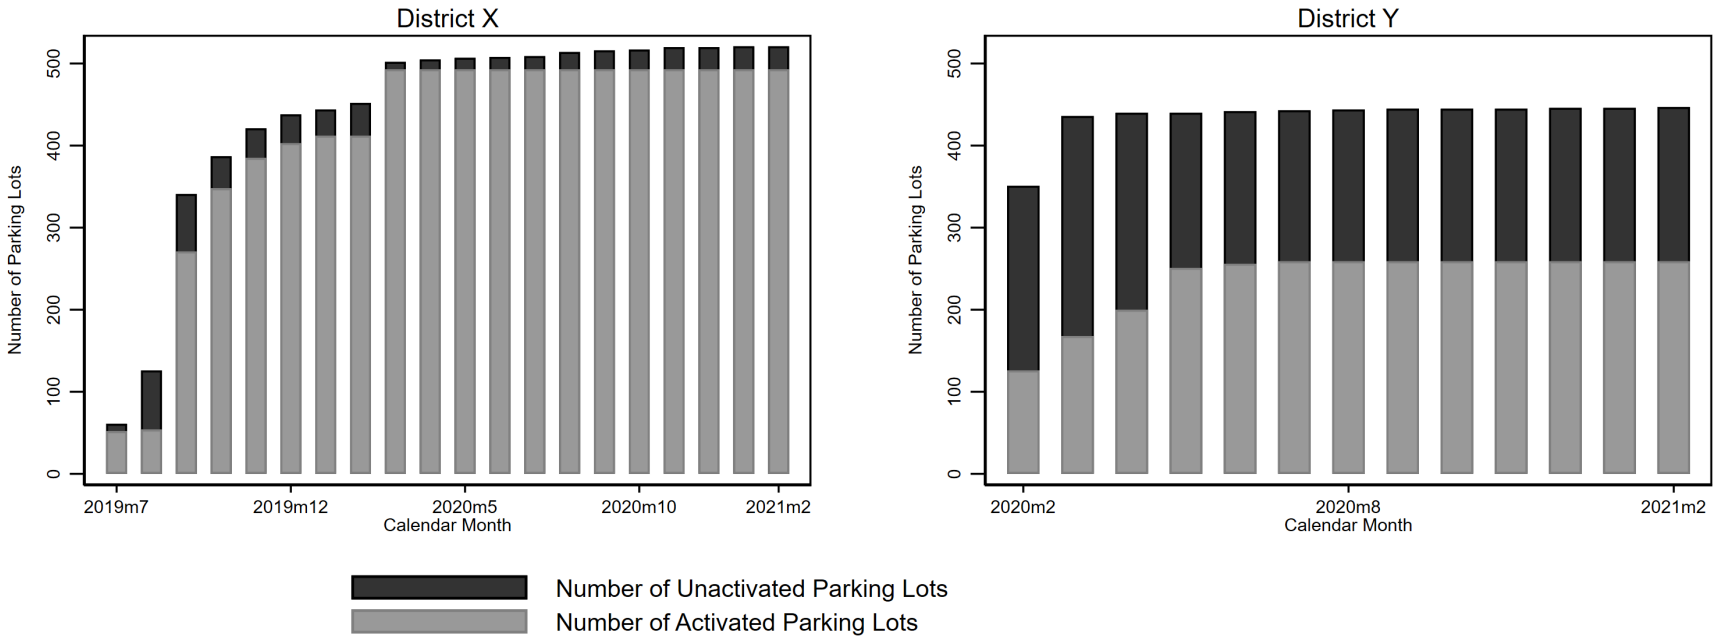
Notes: This figure is a flowchart that describes the process of AI-automated enforcement against illegal parking. Panel A is a snapshot of Company W's dashboard showcasing an illegal parking incident. Once Company W forwards this record to the traffic police, the car owner receives both the photographic evidence in Panel A via an app and the text message displayed in Panel B on their mobile phone. Upon paying the fine through the app, the driver is then issued the electronic ticket (which also serves as a payment receipt) seen in Panel C. Specific details are masked for privacy.

Figure A3. *The Implementation Process in District X*



Notes: This figure illustrates the implementation process in District X. The dataset, denoted by the red horizontal arrow, is utilized to assess the impact of AI-automated law enforcement on illegal parking. The dataset indicated by the black horizontal arrow is employed for additional analyses (as discussed in Section 2.2), which include examining the ripple effects on legal behaviors and conducting a falsification test.

Figure A4. *Staggered Expansion of AI-Automated Enforcement*



Notes: This figure visually depicts the gradual expansion of AI-automated enforcement in Districts X and Y. The height of each bar indicates the number of parking lots equipped with AI-automated illegal parking detection capabilities. The gray portion of each bar represents those lots where the traffic police have activated this capability, meaning photographs from these lots are accepted as evidence for issuing tickets.

Table A1. *Description of Company W's Operation in Districts X and Y*

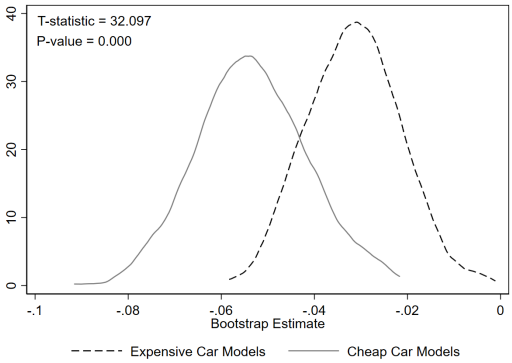
	District	X	Y
Start time to	Record legal parking (Function 1, F1)	January 1, 2019	Febuary 1, 2020
	Detect illegal parking (Function 2, F2)	July 1, 2019	Febuary 1, 2020
Implementation Process		First installing devices for F1 and then upgrading them to accommodate F2.	Installing two separate sets of devices for F1 and F2 simultaneously.
Data Mergability		Data on F1 and F2 could be merged through device IDs.	Data on F1 and F2 could NOT be merged.

Notes: This table details how Districts X and Y have a staggered lot-by-lot adoption of AI-automated enforcement against illegal parking. The last row highlights their differences in data mergeability.

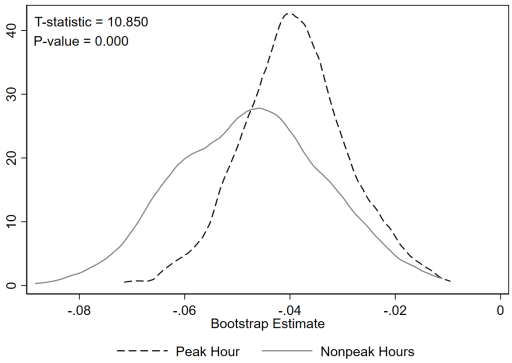
Appendix B Supplementary Results from DID Estimation

Figure B1. *Bootstrap Tests of Distributional Effects*

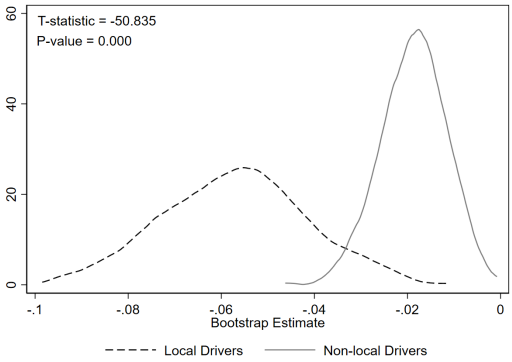
Panel A Expensive vs. Cheap Cars



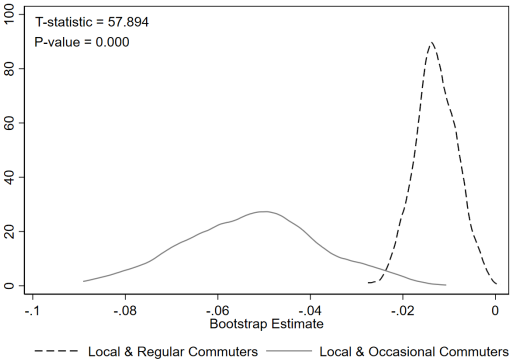
Panel B Peak Hours vs. Non-Peak Hours



Panel C Local vs. Non-Local Drivers



Panel D Regular vs. Occasional Commuters



Notes: This figure presents graphical evidence from bootstrap tests examining distributional effects. Panels A through D represent comparisons between expensive and cheap cars, peak and non-peak hours, local and non-local drivers, and regular and occasional commuters, respectively. T-tests and their corresponding p-values are displayed in the top left corner of each panel.

Table B1. *Summary Statistics of Outcome Variables*

	Full Sample			Treatment Group						Control Group		
				Pre-Treatment			Post-Treatment					
	N	Mean	SD	N	Mean	SD	N	Mean	SD	N	Mean	SD
# of illegal parking	431870	0.890	2.718	23728	1.876	4.548	329808	0.671	1.845	78334	1.515	4.388
Dummy for occurrence of illegal parking	431870	0.277	0.447	23728	0.388	0.487	329808	0.256	0.436	78334	0.329	0.470
# of unique parking episodes	847934	5.482	4.246	146690	5.512	3.972	229629	6.206	4.087	471615	5.120	4.357
Occupied hours	847934	10.233	7.011	146690	9.275	6.354	229629	11.769	6.261	471615	9.783	7.422

Notes: This table reports summary statistics of main variables analyzed in this paper. N and SD stand for observation number and standard deviation, respectively.

Table B2. *Robustness Checks of the DID Model*

	(1)	(2)	(3)	(4)
Outcome Variable		*Log (Illegal Parking)*		*# of Illegal Parking*
Estimator / Model	TWFE OLS	Borusyak et al. (2021)	Cengiz et al. (2019)	Poisson Regression
PostTreat	-0.056***	-0.054*	-0.067***	-0.235**
	(0.018)	(0.032)	(0.025)	(0.092)
Sample	Treatment Group Alone	Full	Full	Full
Parking Spot Fixed Effects	YES	YES	YES	YES
Calendar Date Fixed Effects	YES	YES	YES	YES
Parking Lot FE $\times$ (T, T^2, T^3)	YES			
Date FE $\times$ Parking Lot Attributes	YES	YES		YES
Adjusted R^2 / Pseudo R^2	0.606	/	0.633	0.537
# of Clusters	752	935	7446	966
# of Observations	353536	418563	2781586	431871

Notes: This table presents the results of robustness checks for the DID model. In Column (1), we exclude all control group parking lots. In Columns (2) and (3), we estimate the model using methods proposed by Borusyak et al. (2021) and Cengiz et al. (2019). For Column (4), we employ Poisson regression, drawing upon the model developed by Correia et al. (2019), Correia et al. (2020). To ensure model identification and the convergence of estimation, we had to omit some fixed effects or interaction terms in the last three columns, as indicated in the table. Due to econometric issues (e.g., dropping singletons or stacking data), the number of observations and clusters also differs among columns. Standard errors clustered at the parking lot level are reported in parentheses with ***, **, and * denoting significance at the 1%, 5%, and 10% levels.

Table B3. *Robustness Checks of the Distributional Effects*

Outcome Variable	(1)	(2)	(3)	(4)	(5)	(6)
	Log (Illegal Parking)					
	Min Price		Max Price		Local & Regular	Local & Occasional
	Expensive	Cheap	Expensive	Cheap		
PostTreat	-0.031***	-0.053***	-0.031***	-0.052***	-0.018***	-0.050***
	(0.010)	(0.013)	(0.010)	(0.012)	(0.006)	(0.015)
Sample	Full	Full	Full	Full	Full	Full
Parking Lot Fixed Effects	YES	YES	YES	YES	YES	YES
Calendar Date Fixed Effects	YES	YES	YES	YES	YES	YES
Parking Lot FE $\times$ (T, T^2, T^3)	YES	YES	YES	YES	YES	YES
Date FE $\times$ Parking Lot Attributes	YES	YES	YES	YES	YES	YES
Adjusted R^2	0.529	0.559	0.525	0.562	0.303	0.616
# of Clusters	968	968	968	968	968	968
# of Observations	431870	431870	431870	431870	431870	431870

Notes: This table reports robustness checks of the distributional effects. In Columns (1) – (4), we use the minimum or maximum recommended retail prices (instead of mean) to classify car models as expensive or cheap. In Columns (5) and (6), we extend the time window for defining regular and occasional commuters from 7 to 14 days. Standard errors clustered at the parking lot level are reported in parentheses with ***, **, and * denoting significance at the 1%, 5%, and 10% levels.

Appendix C Results from Regression Discontinuity Design

C.1 Econometric Spcification

Considering that both records of illegal parking and the timing of adoption by traffic police could be accurate to the calendar date, the high-frequency data grants us an opportunity to utilize RDD as an alternative empirical strategy. Different from the classic RDD (see Cattaneo et al. (2019) for comprehensive reviews), the running variable here is the number of days relative to the date of treatment. Hausman and Rapson (2018) have done an excellent job of summarizing the econometric issues when constructing an RDD model with time as the running variable.

Concretely speaking, following the empirical method suggested by Hausman and Rapson (2018) and applied by Li et al. (2020), the RDD model is estimated in a two-step way. First, we estimate the following regression using the full-sample data:

$$Y_{it} = \alpha + \delta^{'} X_{it} + \lambda_i + \theta_t + \varepsilon_{it}\,, \tag{C1}$$

and then use the residual Y_RES_{it} obtained from estimating equation (C1) as the outcome variable to implement the RDD estimation for subsamples within a bandwidth as follows:

$$Y_RES_{it} = \sigma + \beta D_{it} + f(r_{it}) + D_{it} \times f(r_{it}) + \xi_{it}\,. \tag{C2}$$

Y_{it}, α, λ_i, θ_t, X_{it}, and ε_{it} all stand for the same meanings as we have outlined for equation (1). r_{it} is the running variable, and a negative value denotes the number of relative days before the date when the traffic police adopted the picture evidence detected by Company W's monitoring cameras at parking lot i and positive values correspond to the number of dates after that. D_{it} is a dummy variable for treatment status, which equals 1 if r_{it} is no less than 0 and takes the value of 0 otherwise. $f(r_{it})$ is polynomial functions of the running variable and its interaction with D_{it} further allows for distinctive trends of the outcome variable on different sides of the treatment date. ξ_{it} denotes the error term in the second step of estimation.

The main idea behind the two-step procedure is that we first regress the outcome variable against parking lot and date fixed effects (denoted by λ_i and θ_t) along with the two sets of time trends (X_{it}), which could help remove the variation resulted from observables and make the outcome variable less noisy. The relationship between Y_{it} and λ_s, θ_t, X_{it} should hold for the full sample and that is why we use all observations in the first step. After that, we use the less noisy outcome variable (i.e., the residual) for subsequent RDD analyses which could be free from complex specifications, and focus on estimating the discontinuous changes of illegal parking within a narrow time window.

In the baseline setting, we specify $f(r_{it})$ as linear and set the kernel function to be uniform. Since there has not been a consensus on the most reasonable criterion for choices of optimal bandwidth in RDD, we arbitrarily set the bandwidth to be 70 in the baseline setting and experiment with alternative bandwidths in robustness checks, which also cover the optimal bandwidths proposed by Imbens and Kalyanaraman (2012) and Calonico et al. (2014).

Standard errors are clustered at the parking lot level.

Though we have controlled for multiple trends to improve the validity of the DID model in Section 3.2, it is unlikely to exclude all potential omitted variables. As a parallel identification strategy, RDD could help further enhance the credibility of the empirical findings from DID model. Below, we would first display graphic evidence and regression results from RDD under baseline setting, and then experiment with alternative specifications.

C.2 Empirical Results

Figure C1 presents graphic evidence for the RDD estimation, which provides intuitive evidence for the deterrence effect of AI-automated enforcement on illegal parking. The circle stands for the mean within each bin of 5 days and the fitted line along with 95% confidence intervals are plotted with solid and dash lines, respectively. Consistent with the findings from the DID model, we also observe that the adoption of picture evidence detected by AI-automated monitoring cameras could effectively reduce illegal parking.

Column (1) of Table C1 reports the regression result corresponding to graphic evidence in Figure C1. The point estimate implies that the total number of illegal parking will be reduced by 3.6% after the activation of AI-automated enforcement. Compared with the estimated coefficient in DID (6%), this estimate is smaller in magnitude, while still statistically significant at 1% level. To summarize, both DID and RDD results collectively point out that AI-automated law enforcement has been effective in regulating misbehavior like illegal parking.

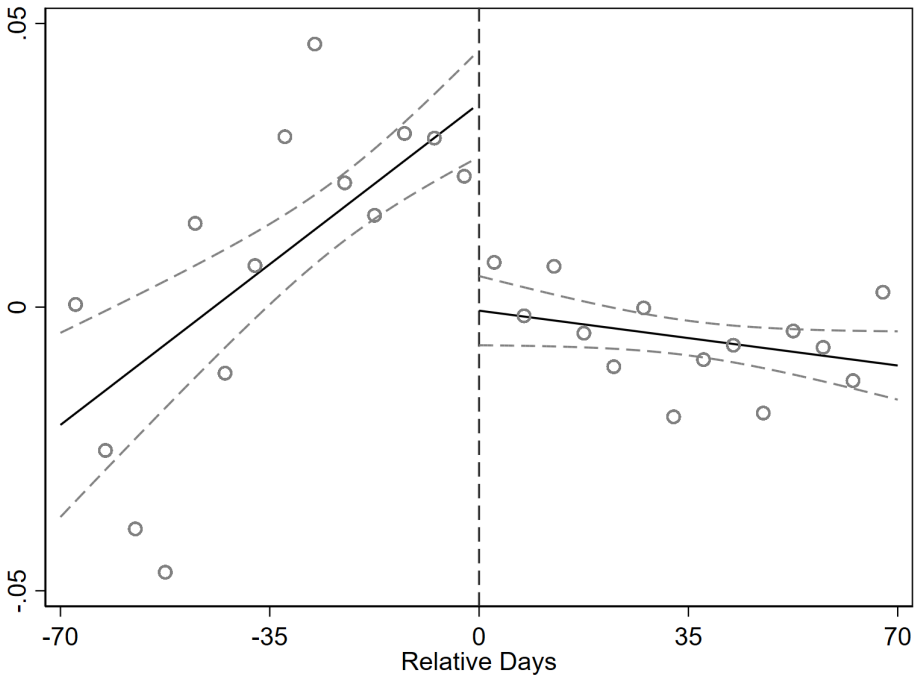
To investigate the robustness of the RDD results, we conduct further analysis on four dimensions. First, the uniform kernel function assigns equal weights to all observations. However, given that the identification of the RDD model highly relies on the jump around the cutoff, some papers (e.g., Lee & Lemieux, 2010) suggest that we should give more weight to the observations closer to the threshold. In Column (2) of Table C1, we switch the kernel function into a triangular one and re-estimate the RDD model. The coefficient implies that the empirical findings are not driven by a specific form of kernel functions.

Second, RDD only utilizes observations within the selected bandwidth to increase the internal validity of the estimation, while sacrificing the data far away from the treatment date. In Column (3) of Table C1, we remove the restriction imposed by bandwidth and use all observations available to re-estimate the model. Additionally, we also add higher-order terms of the running variables to construct the 3-order global polynomial model. The result in Column (3) indicates that the inclusion of more data points would not hurt the main findings and the result may even improve in its magnitude.

Third, since econometricians have not reached a consensus on the criterion for optimal bandwidth in RDD, we have arbitrarily set it to 70 in the baseline model. To investigate whether the findings only hold under some specific bandwidths, we re-estimate Column (1) of Table C1 under different bandwidths which take values between 25 and 115 and change with a step of 1 day. Corresponding point estimates and 95% confidence intervals are plotted in Figure C2. We are glad to find that RDD estimates remain quite robust and exhibit little fluctuations, though some estimates lose statistical significance under bandwidths around 30. It may result from the decreased sample size and consequential loss in statistical power.

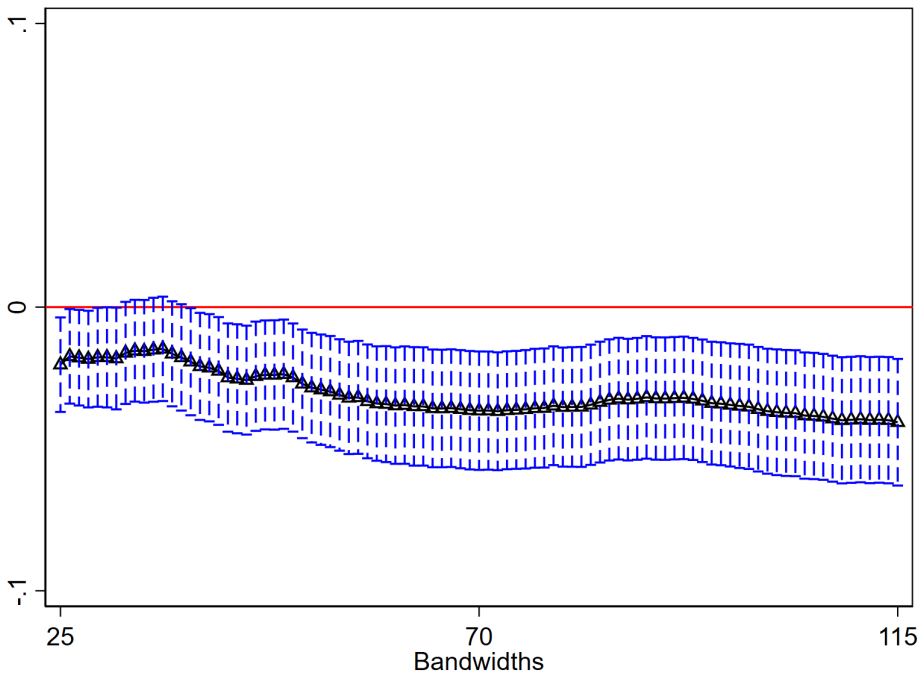
Finally, extant papers differ in choices of standard errors utilized to implement statistical inference. In the baseline model, we use standard errors clustered at the parking lot level. We have further tried heteroskedasticity-robust standard errors, which are reported in square brackets, and find that significance levels change little.

Figure C1. *Graphic Evidence from RDD Estimation*



Notes: This figure presents graphic evidence for the RDD estimation. The circle stands for the mean within each bin of 5 days and the fitted line along with 95% confidence intervals are plotted with solid and dash lines, respectively.

Figure C2. *Estimates under Alternative Bandwidths*



Notes: This figure reports point estimates and 95% confidence intervals under alternative bandwidths. We experiment with bandwidths which take values between 25 and 115 and change with a step of 1 day.

Table C1. *RDD Estimates and Robustness Checks*

	(1)	(2)	(3)
Outcome Variable	*Residuals of Log (Illegal Parking)*		
AfterAdoption	-0.037	-0.028	-0.039
	(0.011)***	(0.009)***	(0.013)***
	[0.006]***	[0.006]***	[0.006]***
Bandwidth	70	70	Global
Kernel Function	Uniform	Triangular	Uniform
Polynomials	LLR	LLR	Cubic
# of Observations	72077	71221	353535

Notes: This table reports RDD estimates and robustness checks. In Column (1), we adopt uniform kernel function and local linear regression (LLR). In Column (2), we switch the kernel function into a triangular one. In Column (3), we use all observations available to re-estimate the model and add higher-order terms of the running variables. Standard errors clustered at the parking lot level are reported in parentheses and heteroskedasticity-robust standard errors are reported in square brackets. ***, **, * denote significance at the 1%, 5%, 10% levels.