From Pumps to Plugs:
Value of Time and Charging Policies

Takeaki Sunada Tianli Xia*

February 15, 2026

Abstract

Designing an electric vehicle (EV) charging infrastructure requires understanding how drivers value charging speed, price, and convenience. We provide the first field-based estimates of drivers' value of time (VOT) in a fast-charging environment, using rich, high-frequency vehicle usage data from Shanghai in 2024. We estimate a random-coefficient discrete-choice model of drivers' charging decisions among the city's commercial chargers. The estimated mean value of time is ¥6.5 per hour (27% of the average wage), implying lifetime charging time costs of about 1.7% of the vehicle price in 2024. Counterfactual analyses show that the city's high-speed charging infrastructure provision since 2020 has reduced the lifetime time cost by 3.9% of vehicle price and accounts for up to 4.6% of new adoptions. We document the relative efficiency of investing on stations and vehicles, as well as possible dynamic inefficiency in the diffusion process. We also show the importance of coordinating investments between station density and speed to achieve an efficient charging infrastructure.

*Sunada: Simon Business School, University of Rochester, tsunada@simon.rochester.edu, Xia: Simon Business School, University of Rochester, tianli.xia@rochester.edu. We thank Luming Chen, Hyuk-soo Kwon, Jim Stock, Yucheng Wang, and seminar participants at Rochester and Salata Institute for helpful comments.

1 Introduction

The transition from gasoline vehicles toward electric vehicles (EVs) is central to global efforts to decarbonize transport and achieve long-term sustainability targets. A key challenge of electric vehicle diffusion is the long charging time, compared to the almost instant fueling for a gasoline vehicle. Charging a 60 kWh battery, the volume that supports up to 7 hours of driving, requires 4-10 hours for a Level II charger and 20 minutes to 1 hour for a Level III charger.[1] Consequently, drivers express the availability, speed, and price of the charging infrastructure as key determinants in the usage and adoption of EVs (Mulcahy et al., 2025; Knittel and Tanaka, 2025). Although charger expansions and technology innovations have accelerated in recent years, their economic impact and welfare implications remain understudied, as little is known about how drivers actually value them.

In this paper, we estimate the value of time (VOT) of EV drivers in the context of vehicle charging decisions, i.e., drivers' willingness to pay to reduce charging time. VOT is the key measure for firm managers and policymakers to design the optimal bundle of multiple policy instruments—such as charging pricing, subsidies, and geographic allocation of high-power stations. Despite the needs, the lack of data has made the direct measurement of VOT challenging. VOT in other contexts (e.g., hourly wage, VOT of gas vehicle drivers) may not accurately represent the drivers' VOT in the charging context, potentially leading to incorrect policy predictions: for instance, drivers may intrinsically dislike waiting, leading to a higher VOT, or drivers who multitask while charging (e.g., eating lunch) may not benefit from faster charging, leading to a lower VOT.

We tackle this challenge by directly measuring EV drivers' VOT using a large field dataset that covers more than 18,000 actual charging decisions involving level-III commercial chargers. We estimate a model of driver decisions on whether and where to charge among stations that offer various charging speeds and prices. Our measure of VOT thus reflects the driver's actual preference for shorter charging times at fast-charging stations, enabling an accurate welfare evaluation of policies aimed at improving charging technologies.

The institution of this paper is Shanghai, China, the city with the largest EV stock in the world, with a population of 25 million, EV stocks of 1.51 million, and 41,000 Level III fast chargers as of 2024. Two features make the setting ideal for studying the driver VOT. First, the widespread adoption of EVs allows us to study a representative set of drivers, while also offering rich variations in observed EV attributes, such as model price and battery range. Second, recent expansion of

[1] Department of Transportation: https://www.transportation.gov/rural/ev/toolkit/ev-basics/charging-speeds.

commercial fast-charging facilities has created a heterogeneous mix of charging speeds and prices, even among fast chargers. These variations in both EV characteristics and charger availabilities make Shanghai an ideal setting to understand the charging decisions of drivers and estimate their VOT.

We leverage two granular and high-frequency datasets on EV charging in the year 2024. The first EV usage dataset tracks the location and battery status among randomly sampled EVs, with detailed information on each charging event, such as GPS coordinates of the charge location, charge duration, and charge volume. We combine the data with a dataset on the universe of commercial fast-charging stations in Shanghai, recording prices and usage at the 5-minute level. The combined data allows us to construct a panel of charging decisions. Importantly, we observe whether a trip leads to charging or not charging (outside options). Given charging, we can identify not only the attributes of the chosen station, but also all nearby stations potentially in the driver's choice set. The data for estimation contains 18,785 trips made by 1,016 drivers, leading to 2,060 fast-charging events across four weeks in 2024.

To identify and estimate the drivers' VOT from their charging decisions — i.e., how driver decisions reflect the prices and charge speed at each station — we adopt a random-coefficient discrete choice model. Drivers decide whether and where to charge their vehicle, for a given set of prices and charge durations available for their expected charge needs. The key parameters of interest are the price and time sensitivity. We control for a rich set of heterogeneity and time-fixed effects, and allow preferences to depend on vehicle attributes (e.g., vehicle prices) and station characteristics (e.g., the number of businesses nearby) to understand the distributions of VOT across demographic groups and across station environments. To account for price endogeneity between stations with potentially unobservable amenity values, we utilize temporal variation in electricity wholesale prices as an instrument.

The results confirm that drivers account for both charging prices and charging time. The average VOT is ¥6.5 ($0.9) per hour, or 27% of the average hourly income in Shanghai. The VOT exhibits great heterogeneity, with the interquartile range of ¥4.2 and a standard deviation of ¥4.9. We find the VOT of the driver increases in the price of the vehicle, our proxy for driver income: while the average VOT among drivers with vehicle whose price is in the bottom quartile is only ¥4.0, that in the top quartile is ¥8.3. We also find that the station amenities substantially affect driver VOT. While stations with no access to nearby businesses and facilities exhibit the driver VOT of ¥9.2, or 42% above the average, the VOT quickly diminishes in the number of nearby businesses

and facilities. This result indicates that drivers concurrently charge their vehicle while performing other tasks, and thus the access to local businesses and facilities is crucial to reduce the drivers' perceived time cost of charging.

Our driver VOT is estimated from driver decisions in the context of commercial fast chargers, making it directly applicable to policy evaluation of public fast-charging infrastructure. We leverage our unified framework in four policy analyses.

First, we quantify the drivers' lifetime time costs from charging. Assuming the charging patterns in the data are representative of the driver, our VOT estimates imply an average annual time cost of ¥574 ($82) from charging, or ¥3,738 ($536) of lifetime cost over 10 years with 10% annual discount.[2] Among drivers with vehicles whose prices are in the top quartile, the annual and lifetime cost amounts to ¥1,035 ($149) and ¥6,740 ($967), respectively. On average, these lifetime charge time costs account for approximately 1.7% of the vehicle price.

We show that the relatively modest time costs primarily reflect the city's dense fast-charging network and drivers' ability to multitask while charging. Counterfactually slowing down all the stations to level-II or eliminating drivers' access to businesses near stations raises lifetime charging costs to 5.6% and 2.6% of vehicle price, respectively, which translates to up to 4.6% and 1.1% lower EV adoptions.[3] These findings indicate that, alongside investments on faster chargers that directly reduce charging duration, strategically locating stations near popular destinations can substantially reduce perceived time costs without costly speed investments.

The second analysis assesses the returns to investments on charging speeds. We evaluate investing in station output and vehicle charging input speeds (known as C-rates), because effective charging speed for a driver is limited by the lower of the two. We find that for current marginal adopters, an investment in charger power has a higher marginal return than an investment in battery input speed. In addition, an investment in faster stations also generates large gains among existing drivers, and overall improved driver surplus can achieve a breakeven 10-year return to investments, even with Shanghai's advanced charger network. Our results speak to policy debates surrounding the Infrastructure Investment and Jobs Act (IIJA) by highlighting public investment in fast-charging infrastructure as a high-return policy lever.[4]

The third analysis quantifies dynamic inefficiency in EV diffusion, that is, ex-post lower-quality

[2] 10 years of ownership corresponds to the total mileage at 233,000km or 145,000 miles.

[3] We assume price elasticity of EV demand of -2.67 estimated in Xing et al. (2021).

[4] Passed in 2022, the IIJA established the National Electric Vehicle Infrastructure (NEVI) program within the Department of Energy, allocating up to $4.155 billion to state governments to cover up to 80 percent of eligible project costs. In January 2025, the incoming Trump administration suspended selected IIJA grant disbursements.

products from early generations are costly to replace, potentially distorting product allocations in subsequent periods. We examine this dynamic inefficiency arising from two channels: *vehicle* side — consumers with high valuation for EVs sort into early-generation vehicles — and *station* side — early stations occupy locations with most charging needs, crowding out later (and more advanced) stations. We study infeasible counterfactuals in which early-generation stations and vehicles are upgraded at 2024 costs, and compare their returns with a feasible upgrade to new stations. We find that inefficiency from the vehicle side is substantial (that is, upgrading early vehicles leads to a higher welfare gain), whereas the station side shows little evidence of dynamic inefficiency. The fact that early stations do not foreclose later-generation stations suggests that subsidies for stations can be a preferred early policy lever over a direct subsidy to vehicles from a long-term point of view.

Our final analysis compares investments in charging speed with investments in station density. We find strong diminishing marginal returns to expanding either dimension in isolation, indicating that a coordinated expansion of both speed and density yields higher overall efficiency. Given the status quo charging infrastructure, investments in station speed deliver larger welfare gains; however, if the average charging speed were 10 kW higher, investments in speed and density would produce similar marginal returns. We also document substantial heterogeneity across income groups: density expansions primarily benefit lower-income drivers, while higher-income drivers with a higher value of time gain more from faster charging. These patterns highlight important equity considerations in the design of charging infrastructure policies.

Related Literature This paper is related to three strands of literature.

The first literature is related to the economics of electric vehicle charging. Early papers have recognized the importance of charging infrastructure in a two-way network setting and highlighted the effectiveness of subsidizing charging stations (Li et al., 2017; Springel, 2021; Li, 2021). More recently, the focus has moved to the intensive margin on how to incentivize charging usage with field experiments (Garg et al., 2024; Bailey et al., 2025). Our study is among the first to examine fast-charging behaviors using field data to quantify the mechanism behind drivers' charging decisions. We find strong sensitivity to charging price, time, and remaining battery volume. Our analysis provides a price elasticity of -3.53 and a duration elasticity of -0.57. We also quantify how a network of charging stations can increase consumers' willingness to pay for a vehicle.

The second literature focuses on the value of time, which has been a main focus underpinning transportation policies since Becker (1965). Most papers estimate consumer VOT based on consumer

preference for trip duration, including studies by Buchholz et al. (2020) and Barwick et al. (2024), and find the VOT close to the hourly wage. In the context of fueling/charging, Houde (2012) and Dorsey et al. (2025) use gasoline fueling decisions to quantify the VOT, while Hoen et al. (2023) estimates using stated preference survey data. Our paper is the first to use field data to quantify the price-time tradeoff based on the actual charge decision of EV drivers. We directly measure the drivers' preference for charging speed, providing a more relevant measure of welfare for policies around improving charging speed.

Lastly, this paper contributes to the literature on the policy design to promote electric vehicles. Many policy debates focus on consumer subsidies (Li et al., 2017), charging infrastructures (Li, 2021), battery and charger access to alleviate range anxiety (Barwick et al., 2025b), electricity pricing (Heid et al., 2024), and innovation policies (Barwick et al., 2025a; Liu, 2025). This paper's main policy focus is to use insights from granular field data combined with micro-founded models to design the optimal bundle of speed, density, and price in fast-charging infrastructure that best supports EV drivers. Such a mix has increasingly become important given the prevalence of fast-charging, especially for commercial vehicles, and drivers in densely populated cities where home-charging is scarce.

The rest of the paper is as follows. Section 2 provides data description and institutional details. Section 3 discusses the model. Section 4 discusses model specification, identification, and estimation strategies. Section 5 reports the main results. Section 6 use our VOT measure to provide policy implications. Section 7 concludes.

2 Research Environment

2.1 Institution and Data Construction

We study Shanghai, a city of 24.9 million residents and 4.15 million households. At the end of 2024, the city had 5.1 million registered passenger vehicles, including 1.51 million EVs— the largest EV stock in any single city worldwide. In 2024 alone, 280,000 new EVs were sold, representing 45% of new passenger vehicle sales, compared to a national average of 40%.[5]

Our primary dataset comes from Shanghai EV Data Center, a government-funded data company that collects and tracks information on all electric vehicles registered in Shanghai. Our dataset

[5]Battery electric vehicles (BEVs) accounted for 71% of the EV stock and 89% of new EV sales in Shanghai in 2024, with the remainder being plug-in hybrids (PHEVs).

consists of four weekly rotating panels of EV driving and charging records in 2024. The weeks are sampled from January, March, August, and October to capture seasonal variation. Each week of panel data contains a random sample of 400 to 600 non-commercial EVs among the population of active vehicles registered in Shanghai, each recording all trips and charging activities by the vehicle during the week.[6] Each trip in the data corresponds to GPS start and end locations during a single spell of power-on, and charging data reports the state of battery volume before and after the charge (in terms of SOC, i.e., percentage of the full battery size), charging duration, and the GPS coordinates of the charge location. We also observe vehicle characteristics, such as brand, model, year, battery size, and manufacturer-suggested retail price (MSRP).

We identify the station used for each charge event by matching the GPS coordinates with a citywide charging station dataset provided by the Shanghai Electric Technology Center. The data covers all commercial level-III fast charging stations operated in 2024, including their GPS coordinates, attributes (such as charger speed and opening date), and prices and occupancy at 5-minute intervals. The granular data allows us to measure charging station availability and prices around each trip. To measure the amenity value of each station, we merge the station data with auxiliary data on the location of points of interest (POI) in the city, which includes all businesses and facilities on Baidu Map in 2024. We also collect data on wholesale electricity prices from the government document.[7] The electricity price is used as an instrument to account for station pricing that is potentially endogenous with unobservable station popularity.

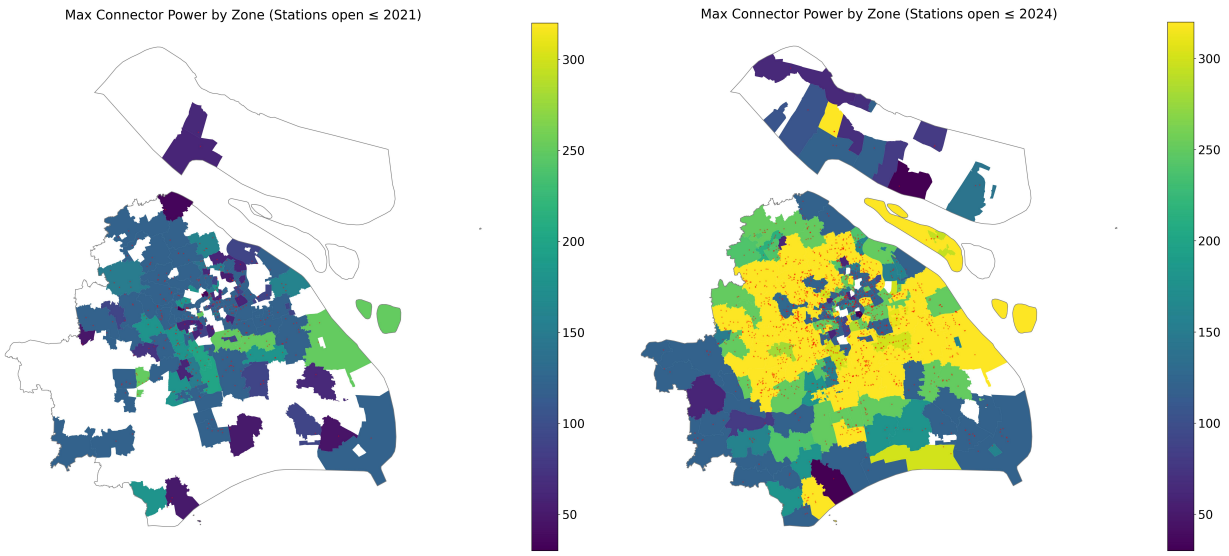
We focus on charging decisions at Level-III public chargers made by drivers without access to a private charger (e.g., at home or workplace). Because our goal is to develop a VOT measure suitable for evaluating policies around public fast-charging infrastructure, we focus on the relevant decision environment. The time–money trade-offs for Level-II or private charging are fundamentally different: drivers with home chargers typically charge overnight and therefore face little or no waiting disutility. In the data, about 27% of drivers fall into our focal category. The excluded drivers include those that only use private chargers (56%) and multihomers (17%).

The final estimation sample consists of 1,016 drivers (vehicles), 18,785 trips, and 2,060 charging events.[8] In the data, 87% of charging events begin within 10 minutes of the preceding trip. Accordingly, we assume that drivers make charging decisions at the trip destination, choosing whether to

[6] Non-commercial drivers exclude taxi and ride-hailing drivers (e.g., DiDi).

[7] As an example, the electricity fee in June 2024 can be found https://ele.aiit.org.cn/static/elefile/9930528202515548717820129887843 6.pdf.

[8] Other dropped trips include ones outside station coverage and ones shorter than 1 km. Appendix A provides the details of our sample construction.



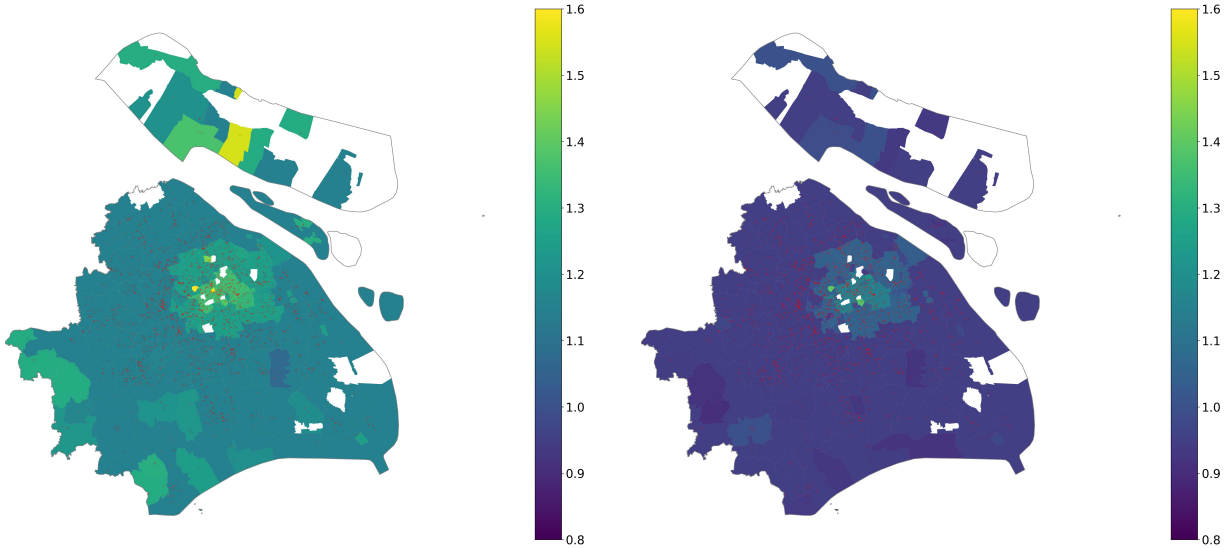
(a) 2021 (b) 2024

Figure 1: Distribution of Fast-Charging Stations and Charger Power

Note: Source: commercial fast-charging dataset. The maps plot the maximum charger power (kW) in each *jiedao* in each year. White areas indicate no fast-charging station available. Lighter areas denote faster chargers. Red dots denote charging stations.

remain plugged in while conducting activities there.[9] This assumption excludes mid-trip detours for charging, which are common for gasoline vehicles but rare for EVs due to long charging times.[10]

To identify the drivers' value of time from their choice of charging station, we also need to construct the set of stations available to the driver during each trip. For our main specification, we assume that the choice set for each trip includes up to five nearest commercial stations located within the min{Trip distance, 2km} radius of the trip destination.[11] We exclude stations that are outside of their hours of operation. Station congestion does not affect our definition of choice set, as the station data reports fully occupied ports less than 1% of the time even during the peak hours. More details on our data construction are available in Appendix A, and robustness checks to alternative sampling assumptions are performed in Appendix D.



(a) Noon (12 PM) (b) 10 PM

Figure 2: Charging Prices across Regions and Time

Note: Source: commercial fast-charging dataset. The maps plot the average charging price (¥/kWh) in each *jiedao* at noon (12 PM) and at night (10 PM). White areas indicate no fast-charging station available at that time. Lighter areas denote higher charging prices. Red dots denote charging stations. Price unit is ¥/kWh.

2.2 Description of Shanghai's Charging Infrastructure

Figure 1 shows the universe of the Level-III charging stations in Shanghai in 2021 (panel a) and 2024 (panel b). We also plot the power of the fastest station at the *jiedao* (street) level (216 in total).[12] White color denotes no stations in that *jiedao*, while lighter colors denote higher power. In 2021, almost half (99) of the *jiedao*s had no charger, while this number dropped to 24 in 2024. More *jiedao* have also gained access to fast-chargers. This across-year comparison shows significant improvement in both station speed and density, whose welfare implications are studied in Section 6.4. The expansion also creates a mix of stations with various speed in the driver choice set, providing key identification variations.

Figure 2 documents spatial and temporal variation in charging prices. Panels (a) and (b) plot the *jiedao*-level distribution of station-average prices at noon and at 10 pm, respectively. On av-

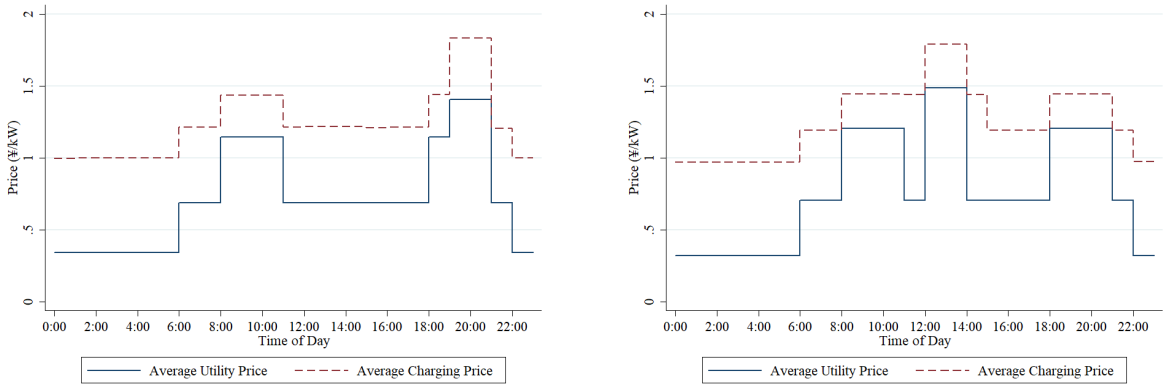
[9] An isomorphic formulation allows charging to occur immediately before the next trip, provided the feasible set of stations is unchanged.

[10] If the observed charging events were performed on route, visiting the observed station would require a median of >57% increase in the trip distance (an average of 416% due to many round trips to the charger) in the data. The uneconomical long detour suggests that these charging events are likely not performed on route.

[11] If a trip ends at a charger, the driver's true destination (to which they walk) is not observed. We take the charger location as the trip destination in this case, with the assumption that the walking distance is short enough so that the choice set would remain the same if defined around the true destination.

[12] *jiedao* is a neighborhood in China and roughly corresponds to a zipcode in the U.S..

erage, daytime prices exceed nighttime prices by 0.7 ¥/kWh, and prices in central Shanghai are consistently higher than those in suburban areas by approximately 0.3 ¥/kWh. Cross-station variation accounts for 31% of the total and reflects persistent heterogeneity in local demand conditions, pricing strategies, and cost structures. Time-series variation accounts for the remaining 60% and is driven by wholesale electricity prices that are common across stations, arising from the utility's peak-load pricing schedule. Figure 3 shows most intertemporal variations in charging prices (dashed lines) reflect those of utility prices (solid lines) both within a day and across seasons.[13]



(a) January (b) August

Figure 3: Charging and Utility Prices by Month

Note: These lines are averages across stations and days during the month. Panel A plots the price for January, and Panel B plots the price for August. Both intraday and seasonal price movements closely track the utility pricing schedule, which we exploit to set prices in Section 3.

2.3 Details on Trips and Drivers' Station Choice

Table 1 reports summary statistics for the estimation sample. Panel (A) presents trip-level characteristics. The average trip lasts 71 minutes, with substantial heterogeneity: 14% of trips are shorter than 15 minutes and 64% are within 60 minutes, while a small share consists of multi-hour intercity trips. Appendix Figure G1 visualizes the geographical distributions of the driving patterns and the associated charger visits in the data. Both the trips and visits are mostly concentrated in city centers, and we also observe trips from and to suburban centers. The mean end-of-trip battery volume is 61% of the vehicle capacity (approximately 40.0 kWh). Drivers initiate charging at an average volume of 38%, add about 33 kWh per session, and spend 43 minutes charging.[14] Appendix

[13] For most stations, pricing schedules are fixed within each calendar quarter (e.g., January–March) and announced well in advance. As a result, unlike gasoline prices, charging prices are highly predictable at the time of the charging decision. We therefore assume that drivers observe prices when making charging choices.

[14] The charge duration data records the actual time spent to fill the vehicle, excluding post-charge parking.

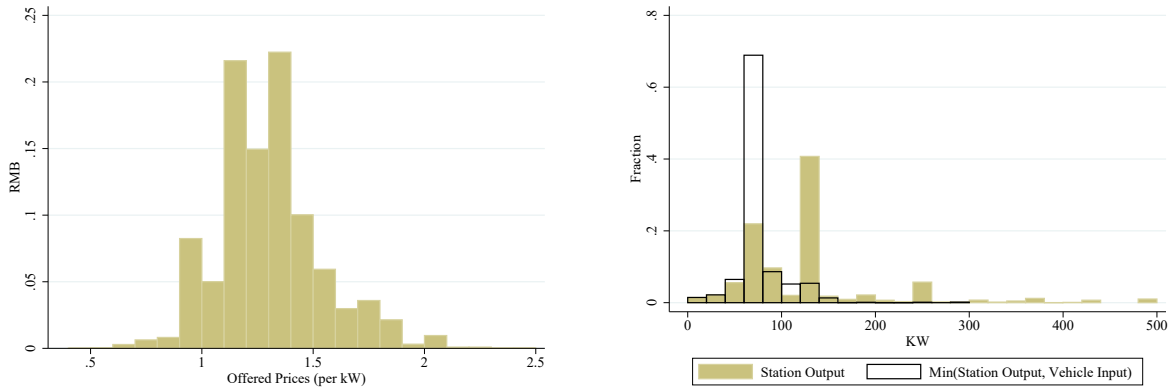
Table 1: Summary Statistics

	Mean	Std. Dev.	Min	Max
(A) Trip-level:				
Trip Duration (Minutes)	70.9	80.5	2	599
End-of-Trip Battery Volume (Overall, %)	61.0	22.5	1	99
(Pre-Charge, %)	38.3	19.8	1	98
Charge Volume (kWh)	32.6	15.4	0.6	93.6
Charge Duration (Minutes)	42.7	19.7	2	179
Price (Available Station, ¥/kWh)	1.30	0.24	0.45	2.89
(Chosen Station, ¥/kWh)	1.21	0.25	0.49	2.42
Electricity Price (¥/kWh)	0.81	0.30	0.32	1.49
Station Power (Available Station, kW)	119.8	80.4	3.5	960
(Chosen Station, kW)	120.8	75.6	20	500
Usable Power (Available Station, kW)	68.9	24.3	3.5	300
(Chosen Station, kW)	68.7	21.9	20	300
(B) Driver-level:				
Vehicle MSRP (¥10,000)	21.1	7.8	10	145.9
Vehicle Battery Size (kWh)	66.7	12.8	37.9	108.8
Maximum Battery Charge Speed (kW)	83.3	47.0	50	315
Weekly Number of Trips	18.4	12.2	1	156
Weekly Trip Duration (Hours, Excluding Charging)	21.7	19.6	0.3	121.0
Weekly Charging Duration (Minutes)	86.1	76.4	2	592
(C) Station-level:				
Points of Interest within 100m	45.4	67.4	0	949
Number of Charging Ports	19.5	19.9	2	194
Number of Trips	18,785			
Number of Drivers	1,016			
Number of Stations	2,027			
Number of Station Visits	2,060			

Note: Source: usage dataset and commercial charging dataset merged for each trip. In panels (A), (B) and (C), the unit of observation is a trip, a driver, and a station, respectively. "Charge Volume", "Charge Duration" and (Chosen Station) rows in panel (A) condition on the trips in which a charge is performed.

Figure G2 plots the distribution of battery volume before and after charging. Charging occurs across a wide range of initial volume levels, but drivers typically charge close to full capacity: the average post-charging volume is 87.8%, and it exceeds 90% in 60% of charging events.

Drivers face an average charging price of ¥1.30 (approximately $0.18) per kWh, with a standard deviation of ¥0.24. The corresponding average wholesale electricity price is ¥0.81 per kWh, implying a station-level operating margin of approximately 39%. Panel (a) of Figure 4 shows a substantial price dispersion, providing ample identifying variations for driver price sensitivity. Drivers select a cheaper station on average (¥1.21 per kWh) and choose the lowest-priced option 53% of the time.



(a) Price

(b) Charging Speed

Figure 4: Distribution of Prices and Charging Speed

Note: The unit of observation is a station-trip pair. In panel (b), colored bars present the station's output speed and uncolored bars present the maximum speed available to the driver in the trip, which is given by the minimum of station output speed and vehicle input speed.

On average, drivers spend ¥39.3 per charging session.

Stations offer an average output speed of 119.8 kW. However, the effective charging speed available to drivers is substantially lower at 68.9 kW because vehicle input speed is the binding constraint for many drivers. In panel (b) of Figure 4, we present the distribution of station output speed (colored bars) and the usable speed for the driver (white bars), showing that the latter is capped at about 80 kW for many charging events. Although charging speed at chosen stations is not markedly higher than average because of the vehicle-side constraints, drivers systematically avoid slow options: in 79% of charging events, they choose the fastest station within the constraints. Appendix Figure G3 compares the price and speed of chosen stations with the best available alternatives in the choice set.

In panel (B), we present the summary statistics of driver-level metrics. We use vehicle MSRP to proxy for the driver's income. The average MSRP is ¥211,000 ($29,600), with the range of ¥100,000 to ¥1,459,000.[15] An average vehicle has a battery size of 66.7 kWh, with the charge input speed of 83.3 kW. Drivers make 18 trips per week on average, totaling 21.7 driving hours, with substantial heterogeneity. Driving intensity varies systematically with income, proxied by vehicle price: drivers in the bottom quartile take about 20 trips and drive 28 hours per week, compared to 16 trips and 17 hours for those in the top quartile. Charging time is also substantial. Figure 5 shows that drivers

[15]The vehicle with the ¥1,459,000 price tag is an outlier in the data, with the second highest MSRP being ¥460,000. We include the outlier for estimation to regulate the model behavior in the high-MSRP range, but exclude it from counterfactual analyses for representativeness.

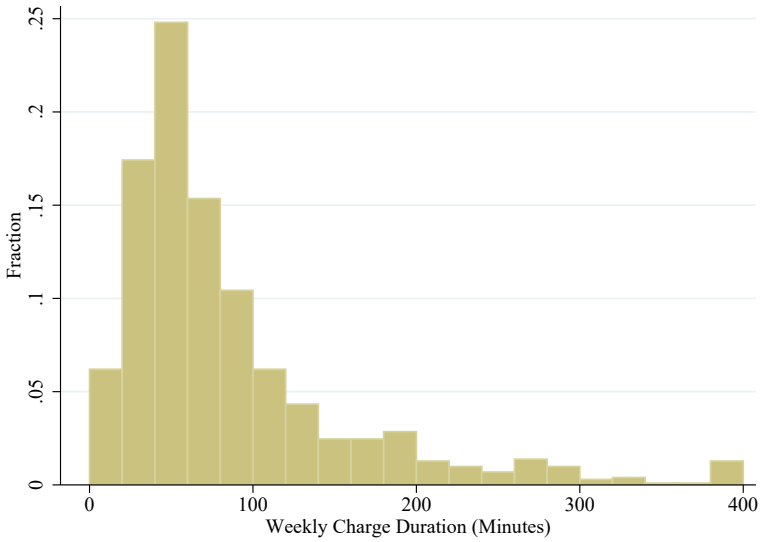


Figure 5: Weekly Charge Duration (Minutes)

Note: The unit of observation is a driver, aggregated across fast charging events during the week.

spend an average of 86 minutes per week charging (median=61; s.d.=76), implying economically meaningful time costs.

Panel (C) summarizes station characteristics used in the demand estimation. We proxy station amenities by the number of points of interest (POIs) within 100 m of the station. Our POI measure includes all businesses and public facilities available in Baidu map, encompassing a broad set of nearby services such as shops, restaurants and hospitals.[16] Stations have, on average, 45.4 POIs within the 100 m radius (s.d.=67.4). We measure station size by the number of chargers, which averages 19.5 per station (s.d.=19.9).

3 Model

To identify the drivers' value of time from their observed choice among charging stations, we specify a random-coefficient, discrete-choice framework. For each trip t, driver i faces a set of available charging stations, $\mathcal{S}_{it}$, around the trip destination, and an outside option of not charging, whose utility is normalized to zero. We specify the utility from charging at station $s \in \mathcal{S}_{it}$ as follows:

$$u_{ist} = \beta_0 - \beta^p_{ist}\bar{q}_{ist}p_{st} - \beta^d_{ist}d_{ist}(\bar{q}_{ist}) + \gamma X_{ist} + f(X'_{ist}) + \xi_{st} + \varepsilon_{ist}, \tag{1}$$

[16] More details on our POI measure are available in Appendix A.

where $\bar{q}_{ist}$ is i's expected charge volume after trip t at station s in kWh, p_{st} is the price per kWh, $d_{ist}(\bar{q}_{ist})$ is the expected charging duration, given as the ratio between the expected charge volume and the charge speed. X_{ist} represents the set of trip-, driver- and station-level covariates. For estimation, we include the remaining battery volume (in %), log vehicle MSRP, vehicle battery size, log station size, and log number of POIs within 100m. We also include day-of-week fixed effects, hour-of-day fixed effects, and the district-level region fixed effects.[17] ξ_{st} represents demand shifters for each station that are unobservable to us and are potentially correlated with p_{st}.

$f(X'_{ist})$ is a flexible function of state variables that might affect the driver's continuation payoff. In practice, charging decisions are inherently dynamic and depend on future charging opportunities. Although we do not explicitly model dynamics, we partly account for them with the $f(X'_{ist})$ term. Specifically, we include a polynomial of the expected battery volume available for the following trip, a variable that likely best captures the driver's future disutility from forgoing charging.[18] The utility from outside option of not charging, u_{i0t}, consists only of this dynamic term and a random utility shock:

$$u_{i0t} = f(X'_{i0t}) + \varepsilon_{i0t}.$$

We allow heterogeneous coefficients for price and charge duration, as follows:

$$\begin{aligned} \log(\beta^p_{ist}) &= \beta^p_0 + \beta^p X^p_{ist} + \sigma^p \nu_i, \nu_i \sim N(0,1), \\ \log(\beta^d_{ist}) &= \beta^d_0 + \beta^d X^d_{ist}, \end{aligned} \tag{2}$$

where X^p_{ist} and X^d_{ist} are the driver-, trip- and station-level characteristics that affect the driver's sensitivity to price and charging duration, respectively. For our empirical application, our main variables of interest are log vehicle MSRP which captures heterogeneity across income groups, and log number of nearby POI which captures utility from concurrent charging ability. We also add a

[17] Walking distance between the station and the trip destination is unobservable (because the car stays at the charger), and thus is excluded from the utility. Although we acknowledge this potentially major omission, it will not bias our VOT estimate if either (1) the trip destination is spatially random within a district, i.e., walking distance is independent of p_{ist} and d_{ist}, or (2) the reason the walking distance is correlated with p_{ist} is because the driver chooses a destination closer to a cheaper station, i.e., the source of the correlation is the driver's response to price. In the latter case, our price coefficient reflects this correlation, which we interpret as the price effect inclusive of the effect through the driver's destination choice.

[18] Formally, this term equals the current battery volume plus the expected charge volume, if the driver chooses to visit a station. Otherwise, it equals the current battery volume. Explicitly accounting for dynamics in the model substantially limits the number of state variables and consumer heterogeneity we can accommodate. Because our goal is to measure driver VOT as flexibly as we can, we opt to account for the dynamics in a reduced-form manner.

lognormally-distributed random coefficient to the price.[19]

We assume that ε_{ist} follows a Type-I Extreme Value distribution. Thus, the probability that each driver chooses a station is given as follows:

$$Pr_{it}(s|\bar{q}_{ist}, X_t, p_t) = \frac{\exp(\delta_{ist})}{\sum_{s' \in \mathcal{S}_{it}} \exp(\delta_{is't})},$$

where $\delta_{ist} = \beta_0 + \beta^p_{ist}\bar{q}_{ist}p_{st} + \beta^d_{ist}d_{ist} + \gamma X_{ist} + f(X'_{ist}) + \xi_{st}$.

Conditional on a station, the driver also chooses the charge volume, q_{ist}. We do not explicitly model how q_{ist} is determined, and instead directly define it as a function of observables and a random consumption shock, ε^q_{ist}, as follows:

$$q_{ist} = \underbrace{\beta^q_0 + \beta^q_p p_{st} + \beta^q_X X_{ist}}_{\bar{q}_{ist}} + \varepsilon^q_{ist}. \tag{3}$$

We allow ε^q_{ist} to be correlated with p_{st}, that is, the price can be endogenous to any demand shocks that affect the average charge volume at the station. However, we assume that ε^q_{ist} is not known to the driver at the time of station choice, and thus the drivers' station choice depends only on $\bar{q}_{ist}$ and not q_{ist}.[20]

In the model, driver i's value of time (VOT) for reduced charging duration at station s on trip t is given by $\beta^d_{ist}/\beta^p_{ist}$. Unlike survey-based proxies, our VOT is identified with the drivers' revealed preference, thus captures their intended use of charging time: VOT may exceed the hourly wage if waiting is intrinsically costly, or fall below it if drivers can multitask during charging. To account for the latter, we allow β^d_{ist} to depend on local amenities, proxied by the number of nearby POIs, and let the data tell whether VOT tends to be lower when nearby amenities are richer.

[19] Adding lognormal duration random coefficients leads to unstable parameter estimates, presumably due to a relatively short panel per driver.

[20] Discrete-continuous framework (Hanemann, 1984) posits that this driver decision on charge quantity can be integrated as part of the utility-maximization framework and thus is used as an identification variation for VOT. We do not opt for this strategy, because we find the available charging speed has almost zero effect on the chosen charge volume, as we show in Section 5.1.

4 Identification and Estimation

4.1 Predicted Charging Volume and Duration

Predicted Charging Volumn Estimation of the station-choice model requires knowledge of $\bar{q}_{ist}$, the planned charging volume, for all the trips, including those that do not involve a charge in the data. We predict its values directly from the data, with an approach similar to Dorsey et al. (2025). Specifically, we estimate expression (3) using the observed charging events. We then use the fitted model to predict the charging volume for all the trips: $\hat{q}_{ist} \equiv E[\bar{q}_{ist}|p_{st}, X_{ist}]$. Because we allow a correlation between the price and the error term, ε^q_{ist}, we use the electricity wholesale price as an instrument for the price.

Predicted Charging Duration Given a planned charging volume $\bar{q}_{ist}$, the expected charging duration is determined by the effective charging speed, i.e., the minimum of station output speed κ_s and vehicle input speed k_i:

$$\bar{d}_{ist} = \frac{\bar{q}_{ist}}{\min\{\kappa_s, k_i\}}. \tag{4}$$

In the data, however, realized charging durations systematically exceed this benchmark because charging speeds taper as the battery approaches full capacity. Since most drivers charge close to full, accounting for tapering is essential for accurately predicting charging duration.

We capture tapering by allowing realized charging time to scale differently across battery ranges:

$$d_{ist}(\bar{q}_{ist}) = \beta_{<70\%}\bar{d}_{ist,<70\%} + \beta_{70\text{–}90\%}\bar{d}_{ist,70\text{–}90\%} + \beta_{>90\%}\bar{d}_{ist,>90\%} + \varepsilon_{ist}. \tag{5}$$

The terms $\bar{d}_{ist,<70\%}$, $\bar{d}_{ist,70\text{–}90\%}$, and $\bar{d}_{ist,>90\%}$ denote the predicted charging time required to traverse each state-of-charge interval in the absence of tapering, computed analogously to expression (4). For example, if a vehicle with a 100 kWh battery charges from 64% to 78%, then $\bar{d}_{ist,<70\%} = 6/\min\{\kappa_s, k_i\}$ and $\bar{d}_{ist,70\text{–}90\%} = 8/\min\{\kappa_s, k_i\}$.[21] Absent tapering, all coefficients equal one. With tapering, we expect $\beta_{<70\%} < \beta_{70\text{–}90\%} < \beta_{>90\%}$, reflecting progressively slower charging near full capacity. We estimate (5) using observed charging events and use the fitted values to predict charging duration for each trip-station pair.

[21] Appendix B formalizes the construction of each component.

4.2 Identification

Identification of our demand model follows the standard identification argument for discrete choice models with micro data. The random coefficient σ^p is identified through the panel structure of the data; we assume each driver's random preference draw, ν_i, remains constant across all of their trips, so that repeated choices reveal individual-specific heterogeneity. In what follows, we discuss two possible threats to identification and how we address them.

Endogenous Price Charging prices may be correlated with unobserved station-level demand shocks, ξ_{st} (i.e., $\mathbb{E}[p_{st}\xi_{st}] \neq 0$). For example, stations at popular destinations may charge higher prices due to strong demand or lower prices to attract traffic. Since the planned charging volume $\bar{q}_{ist}$ depends on prices, both $\bar{q}_{ist}p_{st}$ and the implied charging duration $d_{ist}(\bar{q}_{ist})$ are endogenous.

We instrument $\bar{q}_{ist}p_{st}$ using (i) wholesale electricity prices interacted with station fixed effects and (ii) instruments based on the distance of product attributes proposed in Gandhi and Houde (2019). For the latter, we use the number of stations in the choice set, as well as the sum of the distances in station characteristics (namely, station power, station size, and the number of POIs) between the focal station and the rest of the stations in the choice set. These instruments shift prices while being orthogonal to station-specific demand shocks.

To address the endogeneity of charging duration, we instrument $d_{ist}(\bar{q}_{ist})$ with $1/\min\{\kappa_s, k_i\}$, the inverse of the effective charging speed. Within d_{ist} in expression 4, only the planned charging volume $\bar{q}_{ist}$ is potentially correlated with ξ_{st}; the speed constraint is not. Station output speed is time-invariant and determined at the construction stage, while battery input speed is fixed at the production stage. Any time-invariant correlation between station speed and local geography is absorbed by regional fixed effects and amenity controls. As a result, variations in $1/\min\{\kappa_s, k_i\}$ are plausibly exogenous to trip-specific demand shocks.

Endogenous Choice Set We treat trip destinations, and hence the station choice set, $\mathcal{S}_{it}$, as given, though they are the driver's decision variable in reality. Endogenizing trip formation is beyond the scope of this paper, and our objective is to identify charging choices conditional on realized travel patterns. We therefore discuss under what conditions a potentially endogenous choice set threatens identification and how we address these concerns.

First, the mere fact that the trip destination reflects the drivers' charging needs does not directly cause an estimation bias. Even if drivers choose destinations partly in response to charging needs

(e.g., expected charge volume $\bar{q}_{ist}$), it only induces correlation among observed variables. As long as there is residual variation in $\mathcal{S}_{it}$ conditional on observables, our estimates remain consistent.

Identification concerns arise only if $\mathcal{S}_{it}$ is correlated with unobserved utility components, such as station-level demand shocks ξ_{st} or driver-specific shocks ν_i.[22] To address selection through ξ_{st}, we estimate a robustness specification that includes *jiedao* fixed effects, as well as interactions of day-of-week and hour-of-day fixed effects (168 dummies), in addition to the main controls. This specification restricts identifying variation to comparisons among stations within the same *jiedao* and time window, substantially limiting the scope for endogenous selection on unobserved station attributes.[23] We also include parts of the ξ_{jt} term as residuals from the pricing control function in the utility (as we discuss below), which further absorbs station amenities reflected in prices.

To address selection through ν_i,[24] our baseline specification includes weekend indicators, which capture systematic shifts in drivers' value of time between business and non-business trips and may be correlated with the destination choice. We additionally show robustness that augmenting the model with richer vehicle- and station-level interactions leaves the main estimates unchanged. Taken together, the extensive set of fixed effects and controls implies that remaining variation is plausibly orthogonal to unobserved utility shocks, even in the presence of a potentially endogenous choice set.[25]

4.3 Estimation

Because we estimate the model via maximum likelihood, the instrumental variables are operationalized by the control function approach developed in Petrin and Train (2010). Specifically, for each of the two endogenous terms, we consider the following expression that residualizes the endogeneity:

$$Y_{ist}^k = \gamma_x^k X_{st} + \gamma_z^k Z_{st} + e_{st}^k, \text{ for } k = \{1, 2\} \tag{6}$$

[22] For instance, suppose that a driver anticipating charging needs chooses a destination so that many fast (small d_{ist}) and high amenity (unobservable to us, ξ_{st}) stations are available nearby. In this case, $\mathbb{E}[d_{ist}\xi_{st}|\mathcal{S}_{it}] \neq 0$ even when d_{ist} is orthogonal to ξ_{st} in the population, creating a biased duration coefficient.

[23] Within a *jiedao*, an average of 10.3 stations exist, with a minimum of one and a maximum of 68.

[24] For example, drivers with lower value of time may choose destinations with more amenities (e.g., a mall) and charge concurrently, biasing the estimates because $\mathbb{E}[POI_{st}\nu_i|\mathcal{S}_{it}] \neq 0$.

[25] Intertemporal price variations might prompt drivers to shift their timing of trips, which inflates the number of trips with a visit at a lower-priced time and thus enlarges the price coefficient. Much like the driver's endogenous choice of walking distance discussed in footnote 17, we view these indirect channels as still part of the driver's response to prices, and thus our price coefficient represents the total effect of prices inclusive of them.

where $Y_{ist}^k \in \{\bar{q}_{ist}p_{st}, d_{ist}(\bar{q}_{ist})\}$ and X_{st} and Z_{st} correspond to all the exogenous terms in the utility and instruments, respectively, and e_{st}^k is the error term. We assume the two error terms, e_{st}^1 and e_{st}^2, are jointly normally distributed with ξ_{st}, thus splitting ξ_{st} into its conditional expectation and the residual $\hat{\xi}_{st}$ as follows: $\xi_{st} = \mathbb{E}[\xi_{st}|e_{st}^1, e_{st}^2] + \hat{\xi}_{st}$.

For estimation, we first estimate expression (6) for each endogenous variable via regression. We include the two regression residuals as a linear term in the utility, representing $\mathbb{E}[\xi_{st}|e_{st}^1, e_{st}^2]$, and estimate the distribution of the residual random effect $\hat{\xi}_{st}$ as a nuisance parameter of the model.[26] We adopt a simulated maximum likelihood estimation strategy. We fix 100 simulation draws for the random coefficients and the residual random effect ($\hat{\xi}_{st}$).[27] We search for the optimal sets of parameters Θ to maximize the empirical log-likelihood function:

$$\max_{\Theta} \log \mathcal{L}(\Theta) = \sum_i \log \left(\frac{1}{R} \sum_{r=1}^{R} \prod_{t=1}^{T_i} \prod_{s \in \mathcal{S}_{it}} [Pr_{irt}(s \mid \hat{q}_{ist}, X_{it}; r, \Theta)]^{I(s)} \right),$$

where r corresponds to a draw of random coefficients and the residual random effect, and $I(s) = 1$ if s equals the station selected in the data (including the outside option $s = 0$).

5 Results

5.1 Predicted Charge Volume

We begin by predicting driver i's charging need at opportunity t, $\hat{q}_{ist}$, by regressing realized charging volume on observed station- and trip-level characteristics, along with time and region fixed effects.[28] Table 2 reports the results. Column (1) presents a baseline specification without instrumenting for price. Charging volume is negatively related to price: a one standard-deviation increase in price (¥0.24) reduces charging volume by 1.2 kWh (3.8%). Other station attributes have limited economic impact. Notably, charging speed has essentially no effect on charging volume, indicating that although drivers value speed when choosing stations, they typically charge their desired amount even

[26] With the control function, the utility function to be estimated takes the following form:

$$u_{ist} = \beta_0 - \beta_{ist}^p \bar{q}_{ist} p_{st} - \beta_{ist}^d d_{ist}(\bar{q}_{ist}) + \gamma X_{ist} + f(X'_{ist}) + \underbrace{\beta_{e1} e_{ist}^1 + \beta_{e2} e_{ist}^2}_{\mathbb{E}[\xi_{st}|e_{st}^1, e_{st}^2]} + \hat{\xi}_{st} + \varepsilon_{ist}.$$

[27] A version with 300 draws produces virtually identical estimates for our main parameters of interest.

[28] Station characteristics include price, effective charging speed, log POIs within 100 m, and log number of ports. Trip characteristics include pre-charge battery volume, battery capacity, log vehicle MSRP, and duration of the preceding trip. We also include day-of-week, hour-of-day, and district fixed effects.

at slower stations. In contrast, battery-related variables are strong predictors: lower pre-charge battery volume and larger battery capacity both lead to substantially higher charging volumes. The model explains a large share of the variation in charging behavior (adjusted $R^2 = 0.59$), consistent with the fact that most drivers charge close to full capacity with limited trip-level heterogeneity (Appendix Figure G2).[29] Column (2) instruments for price and is our preferred specification. The price effect becomes larger in magnitude: a one standard-deviation increase in price now reduces charging volume by 1.9 kWh (6.0%). All other coefficients remain quantitatively similar, indicating that price endogeneity primarily biases the price elasticity rather than other dimensions of charging behavior. Column (3) further includes *jiedao*-level fixed effects; while this improves model fit by absorbing regional heterogeneity, it leaves the key coefficients essentially unchanged. In the remainder of the paper, we use the predicted $\hat{q}_{ist}$ from Column (2) to estimate the station-choice model.

	Charge Volume (kWh)		
	No IV	IV	Additional FEs
	(1)	(2)	(3)
Price (¥/kWh)	-4.888 (1.095)	-8.106 (1.116)	-7.591 (1.094)
Usable Charge Speed ($\min\{\kappa_s, k_i\}$ in kW)	0.002 (0.010)	-0.002 (0.010)	-0.001 (0.010)
Log POIs within 100m	0.384 (0.202)	0.483 (0.202)	0.300 (0.225)
Log Station Size	0.495 (0.272)	0.381 (0.271)	0.357 (0.300)
Pre-Charge Battery Volume (%)	-0.486 (0.011)	-0.485 (0.011)	-0.490 (0.011)
Vehicle Battery Size (kWh)	0.405 (0.031)	0.405 (0.031)	0.397 (0.030)
Log MSRP (¥10,000)	4.621 (1.362)	5.035 (1.360)	5.276 (1.340)
Duration of Preceding Trip (Minutes)	-0.004 (0.002)	-0.005 (0.002)	-0.003 (0.002)
Day of Week FE	✓	✓	✓
Hour of Day FE	✓	✓	✓
Region (District) FE	✓	✓	-
Region (*jiedao*) FE	-	-	✓
Price Instrument	-	✓	✓
Adjusted R^2	0.586	0.575	0.625
Sample Size	2,300	2,300	2,300

Table 2: Charge Volume Estimates

Note: Standard errors are reported in parentheses. The sample size is larger than that of the estimation sample because some charging events we drop from the estimation of station choice (e.g., missing information about choice set) can still be used to estimate these regressions.

[29] The duration of the preceding trip has no significant effect on charging volume, suggesting that drivers do not systematically extend trips solely to charge.

	Charge Duration
$\beta_{<70\%}$	0.933 (0.023)
$\beta_{70-90\%}$	1.697 (0.067)
$\beta_{>90\%}$	3.227 (0.146)
Adjusted R^2	0.907
Sample Size	2,308

Table 3: Charge Duration Estimates

Note: The regression excludes intercept. Standard errors are reported in parentheses. A coefficient of one implies that the charge duration equals the charge volume divided by the speed.

5.2 Predicted Charge Duration

In Table 3, we present the estimates of charge duration (expression (5)). We find that the realized charge duration aligns with the theoretical minimum duration up to 70% of the battery capacity, while charging beyond 70% and 90% of the battery capacity requires 70% and 223% extra time, respectively. These results indicate substantial tapering near the vehicle's full capacity. Henceforth, we use the fitted value, $\hat{d}_{ist}$, as the expected charge duration.

5.3 Station Demand Estimates

Table 4 reports the estimated utility parameters from our station choice model. In Column (1), without any instruments, we find an average value of time (VOT) of ¥10.18 per hour, or about 41% of the average wage. Column (2) presents the results from our preferred specification with price instruments. Our estimates of random coefficients imply the average price coefficient of -0.15 with a standard deviation of 0.02, and the average duration coefficient of -0.95 with a standard deviation of 0.62. Compared to Column (1), instruments magnify both price and duration coefficients, indicating the presence of price endogeneity. In terms of elasticities, we find the average own price elasticity of -3.53, with a standard deviation of 2.08.[30] We find the average duration elasticity of -0.57 with a standard deviation of 1.10. These elasticity estimates indicate that drivers respond more to prices than to speed.[31] The predicted average VOT is ¥6.5 per hour, or 27% of the hourly wage, smaller than Column (1).

[30] This estimate aligns with prior findings. For instance, studies report price elasticities of -0.4 for workplace charging at UCSD (Garg et al., 2024), -0.44 for residential charging by Tesla owners in Australia (La Nauze et al., 2024), and -1.6 for home charging in Calgary (Bailey et al., 2025). Evidence from a nationwide UK policy targeting public charging indicates even higher sensitivities, with elasticities ranging from -2 to -2.9 and substantial variation across locations (Bernard et al., 2025).

[31] The full distribution of elasticities is available in Appendix Figure G4.

	Station Choice		
	No IV	IV	Additional FEs
	(1)	(2)	(3)
Linear Parameters			
Pre-Charge Battery Volume (%)	-0.061 (0.014)	-0.136 (0.018)	-0.132 (0.018)
Log Vehicle MSRP (¥10,000)	-1.086 (0.607)	0.138 (0.704)	-0.020 (0.800)
Vehicle Battery Size (kWh)	0.021 (0.013)	0.094 (0.016)	0.090 (0.017)
Log Station Size	0.572 (0.045)	0.594 (0.049)	0.617 (0.057)
Log POIs within 100m	0.072 (0.035)	0.117 (0.039)	0.154 (0.064)
Price sensitivity: $-\exp(\beta_0^p + \beta^p X_{ist}^p + \sigma^p \nu_i)$			
Constant	-1.801 (0.361)	-1.503 (0.201)	-1.570 (0.268)
Log Vehicle MSRP	-0.422 (0.116)	-0.134 (0.065)	-0.133 (0.068)
Weekend Dummy	0.058 (0.090)	0.019 (0.042)	0.020 (0.046)
σ^p	0.301 (0.086)	0.117 (0.019)	0.123 (0.021)
Time sensitivity: $-\exp(\beta_0^d + \beta^d X_{ist}^d)$			
Constant	-5.036 (2.337)	-3.530 (1.596)	-3.456 (2.915)
Log Vehicle MSRP	0.851 (1.128)	0.834 (0.650)	0.749 (0.683)
Log POIs within 100m	-0.184 (0.080)	-0.132 (0.045)	-0.128 (0.040)
Weekend Dummy	0.095 (1.114)	0.044 (0.219)	0.057 (0.191)
Log Station Size	-0.037 (0.087)	-0.038 (0.050)	-0.051 (0.051)
Vehicle Battery Size (kWh)	0.031 (0.023)	0.021 (0.013)	0.023 (0.014)
Dynamic Consideration			
End-of-Period Battery Volume/100	0.112 (0.034)	0.104 (0.039)	0.100 (0.040)
End-of-Period Battery Volume^2/100	-0.115 (0.038)	-0.083 (0.041)	-0.076 (0.047)
End-of-Period Battery Volume^3/10000	0.075 (0.028)	0.030 (0.029)	0.027 (0.035)
Fixed Effects			
Day of Week FE	✓	✓	-
Hour of Day FE	✓	✓	-
Day of Week × Hour of Day FE	-	-	✓
Region (District) FE	✓	✓	-
Region (*jiedao*) FE	-	-	✓
Price Instruments	-	✓	✓
Implied VOT as Hourly Wage			
Mean	41%	27%	27%
Median	25%	20%	20%
Log-Likelihood	-7,857.1	-7,816.4	-7,513.0
Sample Size	18,785	18,785	18,785

Table 4: Utility Parameters

Note: Standard errors are computed with 500 bootstrap simulations. Estimates of first-stage control functions and nuisance parameters are presented in Appendix C. MSRP, manufacturer's suggested retail price; VOT, value of time; FE, fixed effects; POI, point of interest.

For the interactions terms that shift price and time sensitivity, we find that drivers with a more expensive vehicle are less price sensitive and more duration sensitive, consistent with high-income drivers being less sensitive to price and more to time, though the latter effect is not statistically significant. We also find that access to nearby points of interests reduces driver disutility from charging duration, consistent with the hypothesis that multitasking makes charging duration more tolerable. To account for the potential choice-set endogeneity discussed in Section 4.2, we also include additional interaction terms.[32] These coefficients are economically small and statistically insignificant, suggesting that endogenous choice sets will not materially affect our estimates. Notably, weekend dummies have little impact on both price and duration sensitivities, indicating that driver VOT does not vary between business- and leisure-trips. On linear utility terms, we find a higher probability of charging when the driver faces less remaining battery, consistent with the dynamic incentives. We find larger stations and stations with more amenities receive more visits, whereas vehicle characteristics produce unstable coefficients across specifications.

Column (3) augments the baseline specification with interacted day-of-week by hour-of-day fixed effects and *jiedao* fixed effects, substantially expanding the set of controls for potentially endogenous prices and choice sets. Despite the increase in parameters (386 versus 66), the price and duration coefficients remain stable, implying limited bias from unobservable demand shifters and yielding an average VOT of ¥6.5 per hour (27% of the average wage), similar to column (2). We therefore use column (2) as our preferred specification for its parsimonious nature, with additional robustness checks reported in Appendix D.[33] We also present our model fit in Appendix E.

5.4 Drivers' Value of Time (VOT)

Figure 6 presents the distribution of willingness to pay (WTP) for a one-hour reduction in charging time, which is our primary measure of drivers' value of time (VOT). We estimate an average VOT of approximately ¥6.5 (about $0.9), corresponding to roughly 27% of the average hourly wage in Shanghai, with substantial heterogeneity across drivers and locations. The relatively modest estimated VOT is consistent with the transportation literature, which finds that time spent waiting in a vehicle carries a substantially lower opportunity cost than time spent actively driving (Small,

[32] Controls include a weekend dummy in both price and duration coefficients, and log station size and vehicle battery size in the duration coefficient.

[33] Additional specifications include a different definition of choice set, a different definition of POIs and additional interaction terms in random coefficients.

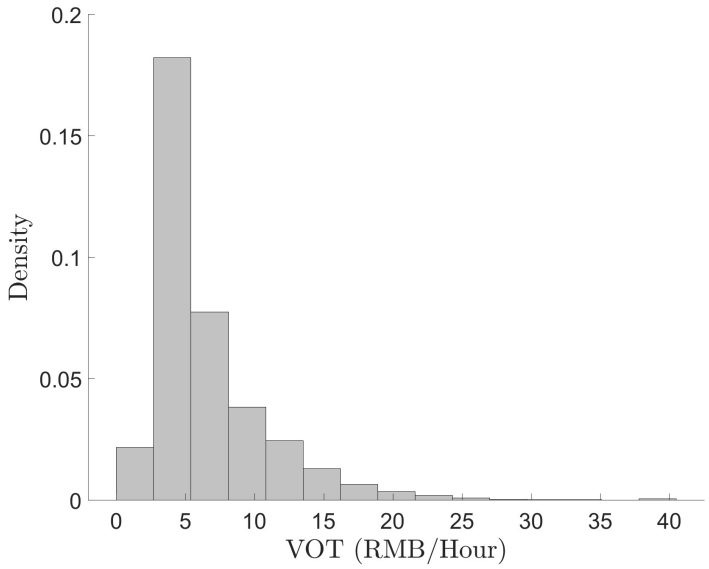


Figure 6: Driver VOT (WTP for One-Hour Reduction in Charge Time)

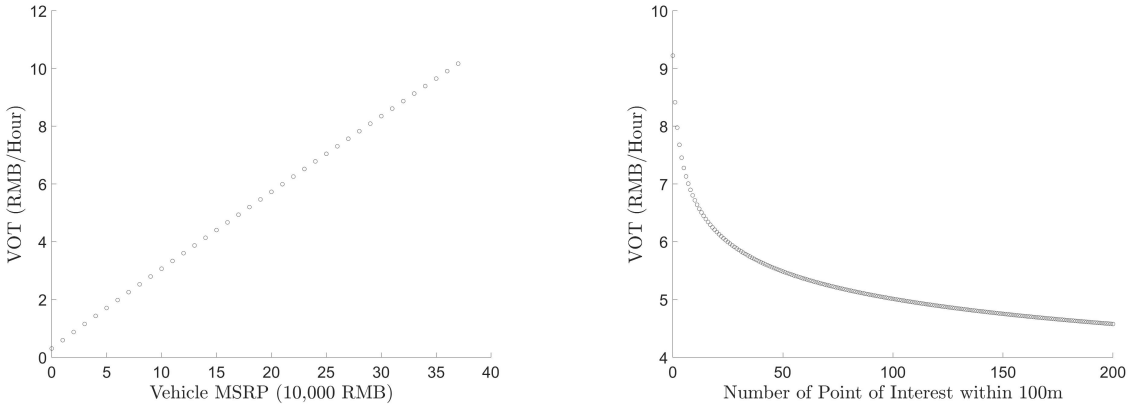
Note: WTP is given by the ratio between the duration and price coefficients. The random part of the price coefficient (ν_i) is simulated with 100 Halton draws. The unit of observation is a trip-station-simulation tuple, and the rightmost bar corresponds to ¥40/hour and higher.

2012).[34]

Figure 7 highlights two key dimensions of heterogeneity in VOT. Panel (a) shows that VOT increases monotonically with vehicle MSRP, indicating that higher-income drivers place a higher value on time. Drivers in the top quartile of the MSRP distribution exhibit an average VOT of approximately ¥8.3, about 28% higher than the population average, whereas drivers in the bottom quartile exhibit a VOT of about ¥4.0, roughly 38% lower. This gradient is consistent with income-related differences in the opportunity cost of time.

Panel (b) shows that VOT varies sharply with station amenities. Drivers exhibit the highest VOT of ¥9.2 (42% above average) at stations with no nearby points of interest, with VOT declining rapidly as amenities become available: even a small number of nearby facilities substantially lowers the perceived cost of charging time. This pattern underscores the importance of multitasking: when charging can be combined with other activities such as shopping or dining, the opportunity cost of waiting is markedly reduced. Because many stations in our data are located in dense commercial areas, this mechanism helps explain the relatively low average VOT in the population. We further explore the welfare implications of concurrent charging opportunities in Section 6.1.

[34]In practice, DOJ recommends using VOT for driving as an hourly wage, while in-vehicle time (not driving) is 0.5 times the hourly wage in this memorandum.



(a) VOT Change in Vehicle MSRP (b) VOT Change in Station Amenity

Figure 7: VOT and Driver - Station Characteristics

Note: The lines are averages across drivers, stations and simulation draws for the random coefficient.

6 Policy Evaluations

Our driver VOT is estimated from driver decisions in the context of commercial fast chargers, making it directly applicable to policy evaluation of public fast-charging infrastructure. We illustrate the applicability of our measure through four policy counterfactuals. First, we document the modest driver lifetime cost of charging and quantify the contribution of the city's charging infrastructure. Second, we assess the return on investments in improving charging speed through increased station power and a higher vehicle input speed. Third, we quantify dynamic inefficiencies in EV adoption, whereby early adopters with a high willingness to pay are more likely to be matched with early, low-quality products. Lastly, we show the importance of coordinated investments in station speed with station density improvements. We also find different income groups value density and speed differently, thus suggesting equity implications for a charging infrastructure investment.

Simulating Driver Time Costs We first discuss how we simulate driver time costs under the factual station network. We assume the weekly trip patterns in the data are representative for each driver and compute the annual time cost as 52 times the predicted weekly cost, abstracting from discounting within the year. To compute the expected weekly time costs consistent with the observed trip patterns, we produce $K = 10,000$ simulated sequences of charging decisions per driver: for each sequence k, starting from the first trip in the data, we simulate the decisions at each trip in which the driver faces the prices and station characteristics observed in the data, and their pre-charge battery volume is determined by their decisions in the preceding trips. We predict

the charge volume and the associated duration using expressions (3) and (5). The annual expected charging time cost of driver i is given as follows:

$$\text{Annual Charging Time Cost}_i = 52 \times \sum_k \frac{1}{K} \sum_t \frac{\beta^d_{is^*t}}{\beta^p_{is^*t}} d_{iks^*t}(\hat{q}_{iks^*t}), \tag{7}$$

where $d_{iks^*t}(\hat{q}_{iks^*t})$ is the time spent charging at chosen station s^* after trip t for simulation sequence k. We use our VOT measure (i.e., $\beta^d_{ist}/\beta^p_{ist}$) to evaluate the associated time costs for the driver. The weekly time cost is obtained by summing the time cost across trips, averaged across simulation sequences. We use the posterior mean of ν_i for each driver i to evaluate the driver's price coefficient, β^p_{ist}.[35]

We also simulate driver decisions under counterfactual scenarios in a similar manner, holding the factual trip patterns fixed.[36] The choice set for each trip is modified as we manipulate station availability and the speed of each station. In principle, any policy changes that affect the utility of charging would also affect the driver's continuation value, which our reduced-form continuation payoff, $f(X'_{ist})$, does not reflect. To address this issue, we adopt Dorsey et al. (2025)'s approach. In the simulation, we manually adjust the outside option value so that the inside good share (i.e., probability of charging in any stations) remains the same between the factual and counterfactuals, thereby predicting the counterfactual utility while ensuring that drivers charge enough to cover their trip needs.[37] We present the step-by-step procedure of simulation implementation in Appendix F.

6.1 Opportunity Cost of Charging Time

We start by documenting the factual distribution of opportunity cost of EV ownership that arises from charging time among the sample of drivers in the data. We then show how the extensive charging infrastructure in the city leads to low driver time costs.

Current Time Costs Figure 8 shows the distribution of the annual opportunity cost of charging time. An average driver incurs a time cost of ¥574 ($82) per year, with a median of ¥399 and a standard deviation of ¥627. Discounted at a 10% annual rate over a 10-year horizon, this

[35] Formally, we use $\mathbb{E}[\nu_i|d_{i1}, ...d_{iT_i}]$ in place of ν_i in expression (2), where $d_{i1}, ..., d_{iT_i}$ is the sequence of charge decisions observed in the data.

[36] Studies on "rebound effect" of energy savings find small elasticity of vehicle miles travelled (VMT) on savings, usually 5%-20%, and therefore our findings are robust to moderate levels of usage changes (Gillingham et al., 2016). Compared to (infeasible) counterfactuals in which drivers also adjust their trip patterns, our simulations produce larger time costs under worse-than-factual counterfactuals, and smaller time costs under better-than-factual counterfactuals.

[37] To avoid the case in which a driver is left with zero battery volume, we also assume that the driver must charge the vehicle if her current battery volume cannot cover the subsequent trip.

corresponds to a lifetime cost of ¥3,738 ($536), or about 1.7% of the average MSRP.[38] Time costs are highly uneven across drivers: those in the top quartile of vehicle MSRP face an average annual and lifetime costs of ¥1,035 and ¥6,740, reflecting their higher value of time and exceeding the population average by roughly 50%.

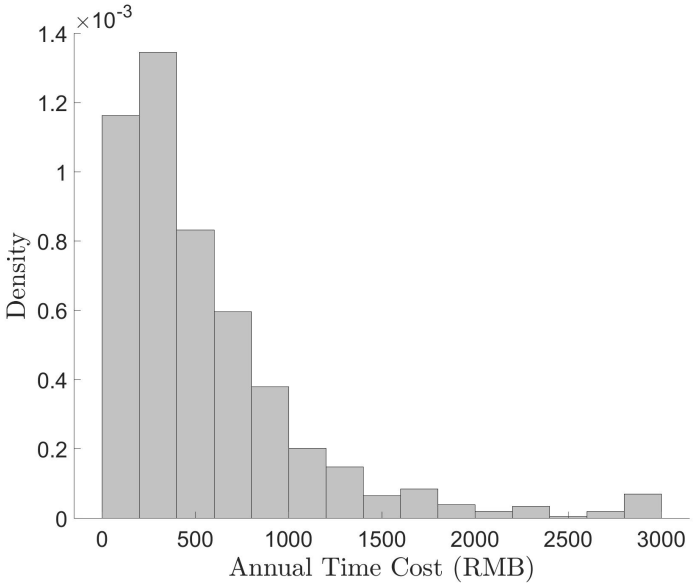


Figure 8: Annual Time Cost of EV Ownership

Note: The unit of observation is a driver, evaluated at the conditional mean of the random coefficient. The rightmost bar corresponds to ¥3,000 and above. The maximum is ¥7,850.

Quantifying the Contribution of Charging Infrastructure We argue that two factors explain the low time cost for drivers in Shanghai: the prevalence of level-III charging stations and the stations' proximity to businesses and facilities. Shanghai has the most extensive Level-III charging network in the world: the city had 41,444 Level-III chargers as of 2024, compared with 3,775 in Los Angeles and 281 in San Francisco. Access to fast-charging stations reduces drivers' time costs by directly reducing the time spent charging. Moreover, most stations are located near major points of interest, allowing drivers to charge their vehicles while concurrently performing other tasks at their destination. We quantify the net contribution of each factor with two counterfactual analyses. In the first counterfactual, we eliminate level-III technology by capping all station speeds at the maximum level-II speed (20 kW). In the second, we shut down concurrent charging opportunities by

[38] Because we abstract from improvements in the charging network over time, these estimates likely represent an upper bound for recent adopters.

setting the number of nearby POIs to zero in the duration coefficient.[39] Together, these counterfactuals isolate how charging environments shape perceived time costs: the former increases charging duration, d_{ikst}, while the latter raises the perceived time cost, β^d_{ist}.

	Factual (1)	All Chargers Level II (2)	Zero Nearby POI (3)
Annual Time Cost			
Average (¥/year)	573.9	1,886.9	888.0
Median (¥/year)	398.7	1,331.1	614.0
Std. Dev.	626.8	2,034.6	968.5
Lifetime Cost (% MSRP)	1.7%	5.6%	2.6%
Δ Implied Vehicle Adoptions (%)	-	[-2.8%, -4.6%]	[-0.6%, -1.1%]

Table 5: Annual and Lifetime Charge Time Costs

Note: These numbers are derived from 10,000 simulated sequences of weekly charging decisions per driver. The implied adoption is defined relative to the factual case. We assume the elasticity at -2.67 when calculating implied adoptions (Xing et al., 2021). The reported changes in adoptions correspond to the minimum (when only single-homers reduce adoption), and the maximum (when both single- and multihomers reduce adoption).

Table 5 reports the results, with column 1 being the factual benchmark. Without level-III chargers (column 2), drivers incur an average annual time cost of ¥1,887, implying a lifetime cost of ¥12,290, or about 5.6% of the vehicle price. Using the price elasticity of EV adoption of -2.67 (from Xing et al. 2021), the extra time cost that is worth 3.9% of the vehicle price would lead to a 10.4% lower vehicle adoption among those who rely on public chargers. In our data, 27% of drivers rely exclusively on public chargers (henceforth referred to as "single-homers") and an additional 17% of drivers use a public charger at least once ("multihomers"). These numbers hence translate to a minimum of 2.8% (if only single-homers reduce adoption) and a maximum of 4.6% (if both single- and multihomers reduce adoption equally) decrease in the overall adoption rate. Therefore, the time cost can be substantial in the absence of fast chargers.[40] When consumers cannot multitask while waiting for charging (column 3), the annual time cost rises to ¥888, corresponding to a lifetime cost of ¥5,784 (2.6% of MSRP) and up to 1.1% lower implied adoption rate. Notably, the gap in the lifetime costs relative to the factual case (0.9% of MSRP) amounts to roughly 24% of the total contribution of the entire level-III network in the city presented in column 2, indicating that

[39] We keep factual POI values elsewhere in the utility function. This counterfactual thus represents a setting in which destinations still provide baseline utility from visiting, but charging cannot be combined with other activities, resulting in higher duration disutility.

[40] Most level-II chargers operate at 7-11 kW rather than the 20 kW assumed herein, so the estimated welfare gains from faster charging are conservative.

reducing the forced waiting is crucial to reduce the drivers' perceived time costs. More broadly, this result highlights that siting fast chargers near amenities can substantially reduce time costs and promote adoptions without costly capital investments in station speed.

6.2 Evaluating Speed Improvements of Stations and Vehicles

We next evaluate the return of increasing charging speed. Because the charge speed for each driver is capped by the minimum of the station output and the battery input speed (as shown in expression (4)), we study improvements on both sides. Specifically, we consider that the social planner spends ¥100 million either to increase the input speed of vehicles sold in 2024, or to increase the output speed of stations built in 2023.[41] Both are costly to improve for technical reasons.[42] According to engineering estimates, adding 1 kW extra power to a charger costs ¥250, and an 1% increase in battery input speed leads to a 0.35%, or approximately ¥186 per unit, increase in battery manufacturing costs.[43]

Surplus from ¥100 mil. Speed Increase	2024 Adopters		All Drivers
	2023 Stations	Vehicle	2023 Stations
	(1)	(2)	(3)
Average per Driver (¥/year)	15.3	6.3	11.1
Median (¥/year)	8.7	2.5	3.1
Std. Dev.	18.0	9.3	18.0
Total across Drivers (¥1,000/year)	[1,158, 1,887]	[478, 779]	[4,928, 8,030]

Table 6: Welfare Gains from Faster Charging Technology

Note: These numbers are derived from 10,000 simulated sequences of weekly charging decisions per driver. A ¥100 million applied uniformly to stations will add 32 kW extra output speed to each charger, and allocating it to vehicles sold in 2024 will add 2.1 kW input speed per vehicle. The reported total gains correspond to the minimum (benefiting only single-homers or 27% of the population) and the maximum (equally benefiting both single- and multihomers or 44% of the population).

[41] We consider investing in stations built in 2023 because many of the 2024 stations were not yet available during the sample periods. Our station data shows 12,375 charging ports in 635 stations built in 2023. The city has a total of 1,643,787 registered vehicles as of 2024, with 280,580 new registrations in 2024. Source: The White book EV Shanghai Data Center Report, 2024.

[42] Faster charging stations are costly because they require powerful, specialized hardware, major electrical grid upgrades (like large transformers), and complex safety systems (isolation links), Improving battery charging speed is costly because it requires advanced hardware to manage significant technical and chemical challenges, primarily related to heat generation and the resulting accelerated battery degradation.

[43] The charger upgrade costs are obtained from Industry report (Guangda Security). The battery charging costs are obtained from Argonne Lab Reports. Note that these costs are to build the product with an already available technology and not the cost of innovation. That is, our definition of investment is to deploy an already feasible technology, and not to innovate something new.

New Adopters In Table 6, we present predicted reduction in annual time costs from these counterfactuals. We first present the gains among the subset of 133 drivers who purchased the vehicle in 2024, a set of newest adopters in the data. In column (1), we present results from uniformly improving the output speed of stations built in 2023 by 32 kW, which costs ¥100 million.[44] The 635 stations built in 2023 account for 33% of all the stations in the data, and 39% of stations in the choice set. We find an average annual gain of ¥15.3 ($2.2) per driver, with a median of ¥8.7 and a standard deviation of ¥18.0. Given the 27% driver population of single-homers and 17% multihomers, these per-driver gains translate to aggregate annual gains of approximately ¥1.2-1.9 million across the 280,580 EV adopters in 2024.

In column (2), we present results from improving the input speed of all the vehicles registered in 2024 by 2.1 kW, equivalent to a ¥100 million investment. We find an average annual gain of ¥6.3 per driver, implying among EV adopters in 2024, the investment produces the total annual gain of approximately ¥0.48 to 0.78 million. Although we study relatively simple uniform investments, the substantial difference in marginal returns suggests that investing on stations is likely a more effective policy lever for increasing the current marginal adopters' welfare, that is, to maximize the speed of EV diffusion.

Existing Adopters While improving battery quality only benefits new adopters, improving station speed extends the benefits to existing drivers. Thus, from the social planner's perspective, the returns from investment for faster stations are substantially larger. In column (3) of Table 6, we consider as outcomes the reduced time costs among all the drivers in the data. We find a lower average driver gain of ¥11.1, with a median of ¥3.1. The lower average return arises because older vehicles with low input speed cannot fully take advantage of the new stations. Nevertheless, given the population of 1,643,787 registered vehicles as of 2024, the total gains across drivers amount to ¥4.9 to 8.0 million, or five times of the gains among the 2024 cohort alone.[45] The estimated gains among drivers alone can yield an approximately 10-year breakeven return to the ¥100 million investment, providing quantitative support for Shanghai's ambitious expansion of super-charging technologies, even from an already high baseline.[46]

[44] We hold the station operators' decisions fixed, and we exclude the gains for station operators. Modeling the optimal supply-side investment and pricing decisions is beyond the scope of the paper.

[45] We assume the same lower and upper bound that corresponds to 27% and 44% of the driver population. In the data, the fraction of single- and multihomers vary little across driver cohorts.

[46] If we assume the 2024 adoption numbers per year for the next 10 years and adopt a yearly discount factor of 10%, the accrued consumer gains reach [¥60 million, ¥97 million], or a 10-year return to investments of [0.60, 0.97]. Note we abstract from many other benefits (e.g., gains among station operators and vehicle manufacturers, and stimulating

6.3 Dynamic Inefficiency in Speed Investments

We next quantify the *dynamic inefficiency*, a phenomenon that is broadly relevant to the diffusion of many technologies. Early adopters—often those with the highest willingness to pay—select into first-generation products that are ex post inferior to later vintages. When technologies such as electric vehicles, charging stations, or solar panels are durable and costly to replace, these early adoption decisions can generate persistent mismatches and create inefficiencies along the diffusion path (Langer and Lemoine, 2022; Chen, 2025).

In the context of EVs, this dynamic inefficiency can arise through two distinct channels. The first operates on the *vehicle side*: consumers who adopt early may remain locked into vehicles with inferior charge speed even as the charger technology advances. The second operates on the *station side*: early charging stations—often built with older technologies—may occupy high-demand locations, crowding out the more advanced chargers in subsequent periods. In this section, we empirically assess the magnitude of these vehicle-side and station-side dynamic inefficiencies. Note that optimal policy designs to address dynamic inefficiency are beyond the scope of this paper because they require a dynamic model of vehicle adoption and charger investment. However, this quantification analysis is still meaningful in offering the first field-based estimates of the dynamic inefficiency in the two-sided EV market.

Vehicle Side We first quantify the inefficiency arising from early EV adopters being locked into vehicles with lower charging speed. We consider an infeasible counterfactual in which the batteries of vehicles registered prior to December 2022 (henceforth referred to as "early adopters") can be upgraded at the cost of 2024. We again consider the social planner's ¥100 million investment. Spread across all the vehicles adopted prior to 2022, this investment raises charging input speed by 0.59 kW per vehicle, a substantially smaller increase (28%) compared to the case of 2024 vehicles studied in Section 6.2 due to the much larger installed base. Despite this dilution, column (1) of Table 7 shows that the upgrade generates an annual surplus gain of ¥1.20 to ¥1.95 million for early adopters or almost three times of the gains among new adopters (column (2), Table 6). Comparing the two scenarios highlights a substantial misallocation: when quality improvements can only be implemented for new vehicles, the resulting inefficiency amounts to 60% of the gains that would be achieved if the same investment could instead be targeted to early adopters.

The dynamic inefficiency reflects the substantial improvements in battery technology over time:

further R&D activities). A formal cost-benefit assessment requires accounting for these other benefits.

early adopters in the data show the average maximum power input of 65.5 kW, compared to 111.5 kW among vehicles registered in 2024. Because of the lower input speed, early adopters spend an average of 92.5 minutes for charging per week, or 38% longer than the 2024 cohort of drivers, even though the early adopters drive only 5% longer.

Station Side We next evaluate the potential inefficiency arising from older charging stations crowding out newer ones at high-demand locations. We consider an infeasible counterfactual in which existing stations can be upgraded at incremental cost only as of 2024, and we allocate ¥100 million to stations built prior to December 2022. This investment raises charging speed by 32 kW across 650 stations with 12,464 charging ports. Column (2) of Table 7 shows that this upgrade generates a total annual welfare gain of ¥5.2 to 8.6 million, which is nearly identical to the return from investing the same amount in the newest stations (column (3), Table 6). The similarity of these returns implies little evidence of meaningful dynamic inefficiency on the infrastructure side.

A plausible explanation is that, although older stations may occupy favorable locations, they do not fully foreclose the entry of nearby newer stations. Consistent with this interpretation, stations built prior to 2022 appear in drivers' choice sets 22.7% of the time, only slightly higher than their 19.3% share of total station counts, suggesting that they do not disproportionately dominate consumers' charging options.

Surplus from ¥100 mil. Speed Increase	$\leq$ 2022 Adopters	All Drivers
	Vehicle	$\leq$2022 Stations
	(1)	(2)
Average per Driver (¥/year)	4.4	11.8
Median (¥/year)	3.1	3.8
Std. Dev.	5.1	24.0
Total across Drivers (¥1,000/year)	[1,198, 1,953]	[5,249, 8,554]

Table 7: Dynamic Inefficiency

Note: These numbers are derived from 10,000 simulated sequences of weekly charging decisions per driver.

Taken together, these results reveal an asymmetry in dynamic inefficiency which appears strong on the vehicle side but very little on the station side. This finding implies another advantage of policies that encourage station constructions in early periods because they incur little inefficiency later on. On the other hand, policies that directly accelerate early vehicle adoption, such as direct subsidies or tax rebates, may lead to unintended mismatches in later periods.

6.4 Welfare and Equity Implications of Station Speed and Density

In our final policy analysis, we compare the welfare gains from improving charging speed with those from expanding station density, with an additional focus on equity considerations in infrastructure design.

Setup We simulate a set of counterfactual environments that jointly vary station density and charging speed. Focusing on stations built on and after January 2022, we construct scenarios in which a subset of stations is removed to reduce density, while the charging speed of the remaining stations is adjusted by -25 to $+25$ kW relative to the factual average of 114.4 kW. For each scenario, we simulate each driver's charging decisions $K = 10{,}000$ times. Because station density benefits drivers beyond the charging speed, we use driver surplus as our main metric, defined as follows:

$$\text{Annual Surplus}_i = 52 \times \sum_k \frac{1}{K} \sum_{t=1}^{T_i} \frac{u_{iks^*t}}{\beta^p_{is^*t}},$$

where u_{iks^*t} is the realized utility from the chosen option s^* in trip t and simulation k, and $\beta^p_{is^*t}$ is the driver's price coefficient evaluated at the posterior mean of ν_i.[47]

In our model, welfare gains from faster charging arise through reduced charging time, while higher station density expands drivers' choice sets with additional options at various prices, speeds, and amenities. The latter also increases the likelihood that a station is available at trip completion, enabling charging in locations where it would otherwise be infeasible or undesirable.

Speed-Density Tradeoff In Figure 9, we present a heat map of welfare changes across the simulated scenarios. The vertical axis represents station density, ranging from only stations built by the end of January 2022 (bottom) to all stations built by the end of 2023 (top), with each step adding stations constructed in a given calendar month; the corresponding number of available stations is shown in brackets. The horizontal axis represents charging speed, varying from -25 kW to $+25$ kW relative to the factual case. On the vertical axis, the heat map exhibits a substantial increase in driver welfare from station expansion. Over the two-year period, an average driver gains approximately ¥354 ($51) per year, most of which is accrued during the year 2023 as more than 600 stations were added to the choice sets. We find a per-station gain of ¥0.54, with a diminishing return: while stations built in the first half of 2022 produced a per-station gain of ¥0.88 per

[47] Realized utility includes ε_{ist}. We set $\xi_{st} = 0$ for all stations and trips, as its distribution is not observed.

driver, those built in the second half of 2023 only produced a per-station gain of ¥0.24. The model predicts decreasing marginal returns because gains from a new choice alternative are decreasing in the number of existing alternatives in the choice set. On the horizontal axis, we find similar diminishing returns from speed investments, as indicated by a flatter indifference curve (the white lines) toward the right.

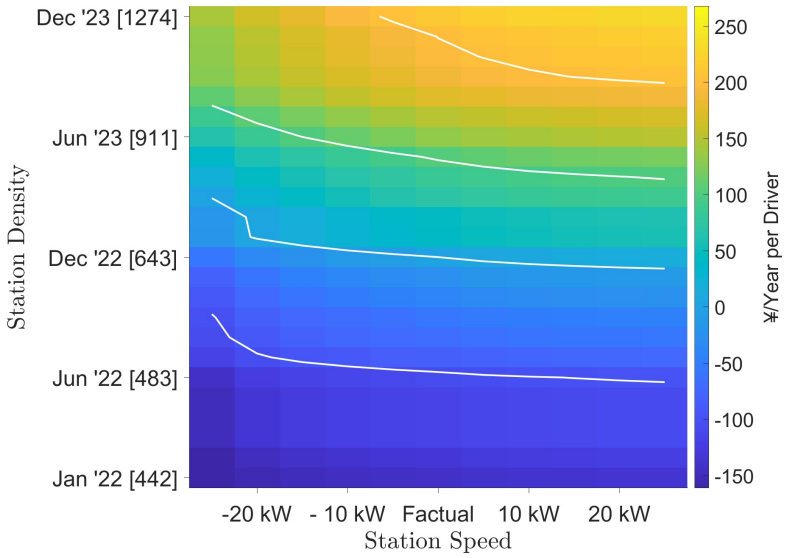


Figure 9: Surplus Change w.r.t. Station Speed and Density

Note: The vertical axis represents station density, and the horizontal axis represents charging speed. Driver surplus is normalized to zero at the center of the figure, corresponding to the set of stations available at the end of 2022 with factual charging speeds. Solid lines denote indifference curves at surplus levels of −100, 0, 100, and 200. The underlying values used to construct the heatmaps are reported in Appendix Table G1.

The rapidly diminishing returns of each policy instrument, as represented in the convex indifference curves, underscore the importance of coordinated investment in both density and speed to maximize driver welfare. For instance, uniformly adding 10 kW to stations built in 2022 produces a gain equivalent of 11.6 new stations (¥8.9) when evaluated at the factual speed (114.4 kW) as of December 2022, whereas the gain would equal 6.0 stations (¥4.6) if the average speed were 10 kW higher.

Based on the engineering cost estimates we use in Section 6.2, the cost of implementing the 10 kW increase is equivalent to building 5.9 stations.[48] Thus, as of 2022, improving speed is the preferred policy at the factual speed, whereas speed and density would produce roughly equal

[48]The investment costs approximately ¥14.9 million, whereas building one station cost ¥2.55 million. The engineering estimates suggest the cost of ¥300 per kW per charger for 2022, and a fixed cost of ¥111,500 per charger.

marginal returns if the average speed were 10 kW higher. We find the same conclusion holds for all the stations built in 2022 and 2023 combined.[49] Thus, while the speed investment has been a more preferred option at the factual speed, a relatively small (10 kW) increase in the available speed would make speed and density investments equally beneficial.

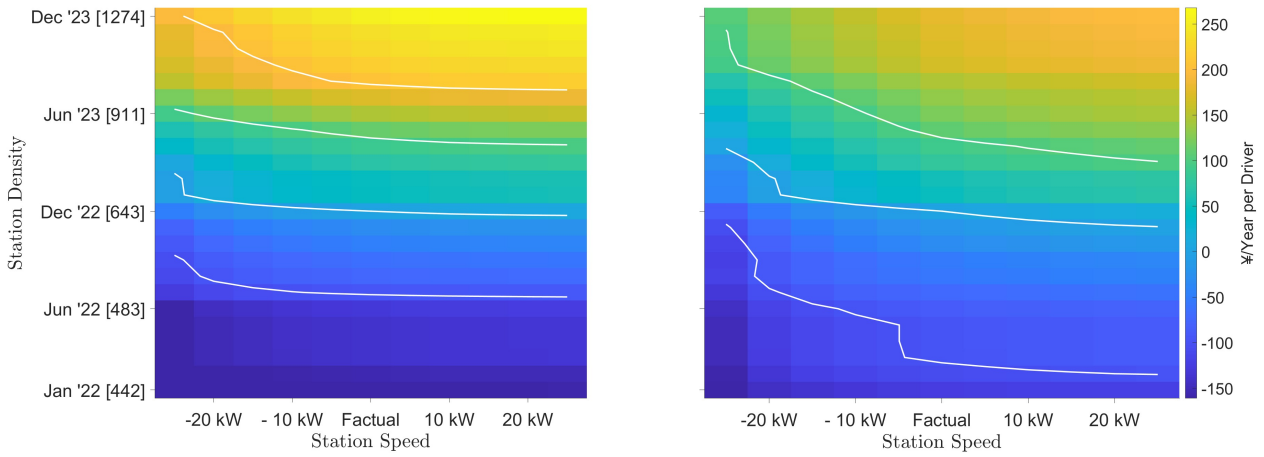
Equity Implications Figure 10 splits the heatmap by vehicle MSRP, comparing drivers in the bottom quartile (panel (a)) and the top quartile (panel (b)) to illustrate heterogeneity across income groups. Along the vertical axis, low-income drivers derive substantially larger gains from increased station density, consistent with their higher vehicle usage: station expansion over the two-year period yields a surplus gain of ¥417 for the bottom quartile, compared with ¥305 for the top quartile. In contrast, high-income drivers benefit relatively more from improvements in charging speed, as seen in the steeper slope of their indifference curves, reflecting their higher value of time. At the factual station density, a uniform 10 kW increase in charging speed generates an average gain equivalent to 16.6 additional stations (¥8.45) for top-quartile drivers, but it is only worth 8.7 stations (¥5.29) for bottom-quartile drivers.

These findings highlight equity implications for charging infrastructure policy: recent expansions in station availability primarily benefit lower-income drivers, whereas investments in charging speed disproportionately favor higher-income drivers and are more likely to accelerate their adoption.

7 Conclusion

In this paper, we estimate the value of time among EV drivers in the context of commercial fast chargers. Using a granular, high-frequency data on vehicle usage and station status, we estimate a model of driver decisions on whether and where to charge. The estimated value of time for an average driver is ¥6.5, or 27% of the average hourly wage, with a significant heterogeneity across driver income groups. The average implied annual time cost during the 10-year period amounts to 1.7% of vehicle price, a modest value thanks to the extensive fast-charging network in the city and the driver's ability to multitask during the charge. Our measure of driver VOT can be used to evaluate various policies around improving charging speed. We find that investing in station power

[49] As of December 2023, improving all the stations by 10 kW will lead to a gain equivalent to 31.5 new stations at the factual speed, and a gain equivalent to 19.0 stations at the 10 kW higher speed. The cost of improving all stations' speed by 10 kW is equivalent to building 16.6 new stations, evaluated at the cost of ¥250 per kW per charger and a fixed cost of ¥111,500 per charger.



(a) MSRP Bottom Quartile

(b) MSRP Top Quartile

Figure 10: Surplus Change by Vehicle MSRP

Note: Surplus at the center of each heatmap (Dec 2022 at factual speed) is normalized to zero. The lines are indifference curves plotted at the surplus of -100, 0, 100 and 200 from the bottom to the top. The number of available stations in the vertical axis is presented in brackets. The underlying values used to construct the heatmaps are reported in Appendix Tables G2 and G3.

is a preferred policy lever over improving vehicle input speed, both in terms of an increased EV diffusion rate and the benefits to exiting drivers. We also find substantial dynamic inefficiency in vehicle adoption (i.e., high WTP consumers are matched with ex-post low-quality vehicles), whereas station from early generations do not seem crowding out newer-generation stations, a result that further justifies subsidies to stations as early policy lever over subsidies to vehicles. Finally, we compare the welfare implications between station speed and density investments and show that while low-income drivers primarily benefit from station density, high-income drivers comparatively benefit more from station speed.

This paper is a first attempt to obtain a realistic measure of driver VOT, i.e., their willingness to pay for reduced charging time. Because our measure is estimated from the variation of drivers' decisions around level-III fast chargers, the estimated VOT is most useful for evaluation of policies around commercial fast chargers. This VOT however is likely context specific. For instance, charging decisions around level-II home chargers (whether they keep the vehicle plugged overnight or charge at other times) likely involve quite different types of trade-offs, thus our measure likely does not apply to, e.g., studying optimal subsidy to facilitate adoption of home chargers. Comprehensive policy evaluations around charging speed require more measures from various contexts, to be left for a future study.

References

Bailey, Megan R., David P. Brown, Blake Shaffer, and Frank A. Wolak, "Show Me the Money! A Field Experiment on Electric Vehicle Charge Timing," *American Economic Journal: Economic Policy*, May 2025, *17* (2), 259–284.

Barwick, Panle Jia, Hyuk-Soo Kwon, Shanjun Li, and Nahim B. Zahur, "Drive Down the Cost: Learning by Doing and Government Policies in the Global EV Battery Industry," January 2025.

_ , Shanjun Li, Andrew Waxman, Jing Wu, and Tianli Xia, "Efficiency and Equity Impacts of Urban Transportation Policies with Equilibrium Sorting," *American Economic Review*, October 2024, *114* (10), 3161–3205.

Barwick, Panle, Shanjun Li, and Tianli Xia, "Range Anxiety," *Working paper*, 2025.

Becker, Gary S., "A Theory of the Allocation of Time," *The Economic Journal*, September 1965, *75* (299), 493–517.

Bernard, Louise, Andy Hackett, Robert Metcalfe, Luca Panzone, and Andrew Schein, "The Impact of Dynamic Prices on Electric Vehicle Public Charging Demand: Evidence from a Nationwide Natural Field Experiment," Technical Report, Centre for Net Zero May 2025. Working Paper.

Buchholz, Nicholas, Laura Doval, Jakub Kastl, Filip Matejka, and Tobias Salz, "The Value of Time: Evidence From AuctionedCab Rides," 2020.

Chen, Luming, "The Dynamic Efficiency of Policy Uncertainty: Evidence from the Wind Industry," January 2025.

Dorsey, Jackson, Ashley Langer, and Shaun McRae, "Fueling Alternatives: Gas Station Choice and the Implications for Electric Charging," *American Economic Journal: Economic Policy*, 2025.

Gandhi, Amit and Jean-François Houde, "Measuring Substitution Patterns in Differentiated Product Industries," Working Paper 2019.

Garg, Teevrat, Ryan Hanna, Jeff Myers, Sebastian Tebbe, and David G. Victor, "Electric Vehicle Charging at the Workplace: Experimental Evidence on Incentives and Environmental Nudges," October 2024.

Gillingham, Kenneth, David Rapson, and Gernot Wagner, "The Rebound Effect and Energy Efficiency Policy," *Review of Environmental Economics and Policy*, January 2016, *10* (1), 68–88. Publisher: The University of Chicago Press.

Hanemann, W Michael, "Discrete/Continuous Models of Consumer Demand," *Econometrica*, May 1984, *52* (3), 541–561.

Heid, Pascal, Kevin Remmy, and Mathias Reynaert, "Equilibrium Effects in Complementary Markets: Electric Vehicle Adoption and Electricity Pricing," *CRC TR 224 Discussion Paper Series*, December 2024. Number: crctr224_2024_615 Publisher: University of Bonn and University of Mannheim, Germany.

Hoen, Fredrik Solvi, María Díez-Gutiérrez, Sahar Babri, Stephane Hess, and Trude Tørset, "Charging Electric Vehicles on Long Trips and the Willingness to Pay to Reduce Waiting for Charging. Stated Preference Survey in Norway," January 2023.

Houde, Jean-François, "Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline," *American Economic Review*, August 2012, *102* (5), 2147–2182.

Knittel, Christopher R. and Shinsuke Tanaka, "Chapter 19: Challenges to expanding EV adoption and policy responses," in "in" June 2025. Section: Handbook on Electricity Regulation.

La Nauze, Andrea, Lana Friesen, Kai Li Lim, Flavio M. Menezes, Lionel Page, Thara Philip, and Jake Whitehead, "Can Electric Vehicles Aid the Renewable Transition? Evidence from a Field Experiment Incentivising Midday Charging," CESifo Working Paper 11386, CESifo, Munich 2024.

Langer, Ashley and Derek Lemoine, "Designing Dynamic Subsidies to Spur Adoption of New Technologies," *Journal of the Association of Environmental and Resource Economists*, November 2022, *9* (6), 1197–1234. Publisher: The University of Chicago Press.

Li, Jing, "Compatibility and Investment in the U.S. Electric Vehicle Market," *R&R AER*, 2021, p. 53.

Li, Shanjun, Lang Tong, Jianwei Xing, and Yiyi Zhou, "The Market for Electric Vehicles: In-

direct Network Effects and Policy Design," *Journal of the Association of Environmental and Resource Economists*, March 2017, *4* (1), 89–133. Publisher: The University of Chicago Press.

Liu, Zikun, "Policy Design for Sustainable Innovation and Adoption: A Structural Analysis of the Electric Vehicle Market," 2025.

Mulcahy, Rory, David Fleischman, Margarietha J de Villiers Scheepers, and Nicolas Reichenbach, "Proud to go electric: Overcoming perceived functional barriers to EV adoption through congruent messaging frames," *Transportation Research Part F: Traffic Psychology and Behaviour*, October 2025, *114*, 1–15.

Petrin, Amir and Kenneth Train, "A Control Function Approach to Endogeneity in Consumer Choice Models," Journal of Marketing Research 2010.

Small, Kenneth A., "Valuation of travel time," *Economics of Transportation*, December 2012, *1* (1), 2–14.

Springel, Katalin, "Network Externality and Subsidy Structure in Two-Sided Markets: Evidence from Electric Vehicle Incentives," *American Economic Journal: Economic Policy*, November 2021, *13* (4), 393–432.

Xing, Jianwei, Benjamin Leard, and Shanjun Li, "What does an electric vehicle replace?," *Journal of Environmental Economics and Management*, May 2021, *107*, 102432.

Web Appendix

A Data Construction

Of the raw data that includes approximately 4,000 drivers, we first drop drivers who have access to a private charger. We classify a driver as having access to a private charger if they either (i) charge their vehicle at home overnight at an average charging speed below 20 kW, or (ii) charge the vehicle at their workplace at an average charging speed below 20 kW. Home and workplace locations are inferred as the vehicle's most frequent parking locations during nighttime hours (12 a.m.-6 a.m.) and weekday working hours (9-11 a.m. and 2-4 p.m.), respectively. Besides home and work chargers, we also find that GPS coordinates of some charging events do not match with the station data, suggesting some charging events at a private charger. We assume charging is at a commercial station if (1) the matched station is the closest among all the commercial stations, and it is within 1km of the coordinates at which the charging event took place, (2) a level III charger is used for the charge, (3) the coordinate is not the driver's home or workplace, and (4) the recorded charge duration is less than min{four hours, 10 $\times$ what a normal charging event would take for the observed charge volume}. We classify any charging events as performed at a private charger if any of these conditions are not met.

After identifying charging events at commercial chargers, we drop from the estimation sample the set of drivers who have access to private chargers, and those who never visited a commercial station during the sample period. We further drop short trips with less than 1km of traveled distance, trips with missing coordinates, and trips for which the predicted charge volume (described in Section 5.1) is below zero. In the estimation, we focus on trips for which at least one charging station is within the choice set, because otherwise, the drivers cannot charge mechanically. These sample selection criteria leave 1016 drivers making 18,785 trips for our estimation, of which 2,060 (11%) include a visit to a commercial charging station. Approximately 76% of the trips include five stations in the choice set, and about 5% to 7% of the trips include 1 - 4 charging stations each. On average, the closest charging station is 550m away from the trip destination.[50]

Some of the stations in the station characteristics data report an unusually low charge speed. Because most stations witness at least one visit during the data period, we impute the station's output power, κ_s, with q_{ist}/d_{ist}, i.e., per-kW time spent at the observed charging event (if multiple

[50] For trips that visited a station, the closest station is the visited station and mechanically has the distance of zero.

visits are observed, their average). Because this measure does not account for tapering (discussed in Section 4.1), the imputed power is a conservative measure of the actual output power. These instances consist of 4.6 % of the data.

Our measure of points of interest (POI) is obtained from Baidu map and includes all the businesses and public facilities available in the city. The data includes a total of 875,978 items. The most commonly observed category is Enterprise, showing up about 120,000 times, followed by Chinese Restaurants (66,000 times). Each business counts as one unit, thus a mall is considered as a collection of many POIs.

B Definition of $\bar{d}_{ist,\%}$ in Expression (5)

Formally, each term in expression (5) is given as follows (SOC: volume in % of battery size);

$$\bar{d}_{ist,<70\%} = \frac{\max\{0, (\min\{0.7, \text{Charge end SOC}\} - \text{Charge start SOC}) \times \text{Battery size}\}}{\min\{\text{Station output speed, Vehicle input speed}\}},$$

$$\bar{d}_{ist,70-90\%} = \frac{\max\{0, (\min\{0.9, \text{Charge end SOC}\} - \max\{0.7, \text{Charge start SOC}\}) \times \text{Battery size}\}}{\min\{\text{Station output speed, Vehicle input speed}\}},$$

$$\bar{d}_{ist,>90\%} = \frac{\max\{0, (\text{Charge end SOC} - \max\{0.9, \text{Charge start SOC}\}) \times \text{Battery size}\}}{\min\{\text{Station output speed, Vehicle input speed}\}}.$$

C Estimates of Nuisance Parameters

Table C1 shows the first-stage estimates of the control function (expression (6)). Column 1 corresponds to the price and column 2 corresponds to the duration. Our primary instruments, namely electricity price for the price and the inverse speed constraint for the duration, are both strong predictor of the dependent variables. Table C2 presents the estimates of nuisance parameters we included in the station choice model. The first two rows correspond to the coefficients of the residuals of the control function ($\mathbb{E}[\xi_{st}|e^1_{ist}, e^2_{ist}]$), and the last row corresponds to the standard deviation of the residual random effect, $\hat{\xi}_{st}$. We find both residuals have a meaningful impact on the estimates, suggesting the presence of endogeneity, and the remaining nuisance random effect is quite small.

D Robustness of Utility Parameter Estimates

In this section, we provide estimates of utility parameters (corresponding to ones in Table 4) utilizing alternative specifications, thus showing robustness of our estimation strategy. In column 1 of Table

	First Stage Control Function	
	$\bar{q}_{ist}p_{st}$	$d_{ist}(\bar{q}_{ist})$
	(1)	(2)
Linear Parameters		
Pre-Charge Battery Volume (%)	-0.629 (0.005)	-0.012 (0.0002)
Log Vehicle MSRP (¥10,000)	6.894 (0.278)	0.083 (0.014)
Vehicle Battery Size (kWh)	0.650 (0.029)	-0.022 (0.001)
Log Station Size	-0.294 (0.069)	0.007 (0.002)
Log POIs within 100m	0.337 (0.059)	0.004 (0.002)
Dynamic Consideration		
End-of-Period Battery Volume/100	-0.278 (0.100)	0.025 (0.006)
End-of-Period Battery Volume^2/100	0.560 (0.160)	0.017 (0.009)
End-of-Period Battery Volume^3/10000	-0.510 (0.084)	-0.011 (0.004)
Instruments		
Electricity Price (Average)	6.271 (0.224)	-
Sum Δ Connector Power (10^{-3})	-0.323 (0.531)	-
Sum Δ Station Size (10^{-3})	-0.766 (0.629)	-
Sum Δ POIs within 100m (10^{-3})	0.533 (0.155)	-
# Other Stations in Choice Set	-0.055 (0.036)	-
Inverse of Speed Constraints ($1/\min\{\kappa_s, k_i\}$)	-	28.345 (1.585)
Fixed Effects		
Day of Week FE	✓	✓
Hour of Day FE	✓	✓
Region (District) FE	✓	✓

Table C1: Control Function

Note: Standard errors are computed with 500 bootstrap simulations. Sum Δ terms are the sum of the differences in the attribute between the focal station and all the other stations in the choice set. Electricity price is interacted with station fixed effects, and their average is reported in the table. MSRP, manufacturer's suggested retail price; VOT, value of time; FE, fixed effects; POI, point of interest.

	Nuisance Parameters
e^1_{ist}	0.141 (0.018)
e^2_{ist}	1.848 (0.751)
$\sigma_{\hat{\xi}_{st}}$	0.1e-6 (0.019)
Adjusted R^2	0.907
Sample Size	2,308

Table C2: Nuisance Parameters

Note: Standard errors are computed with 500 bootstrap simulations.

D1, we present a specification in which we include 10 nearest stations (compared to 5 in the main specification) from the trip destination in the choice set. The specification produces the average price coefficient of -0.15 with a standard deviation of 0.02, and the average duration coefficient of -0.92 with a standard deviation of 0.64. The average implied VOT is ¥6.4 or 26% of the average wage, with a median of 5.0 and a standard deviation of 4.2. In column (2) of Table D1, we present a specification in which we define the station amenity value by the number of points of interest within 200m, as compared to 100m in the main specification. On average, stations have access to 125.8 points of interest, with a standard deviation of 157.1. The minimum and maximum POIs is 0 and 1,407, respectively. This specification produces the average price coefficient of -0.14 with a standard deviation of 0.02, and the average duration coefficient of -0.89 with a standard deviation of 0.58. The average implied VOT is ¥6.9, or 28% of the average wage, with a median of 5.2 and a standard deviation of 5.3. How VOT changes in the station amenity values is also similar to the main specification. In Figure D1, we reproduce panel (b) of Figure 7 with the current specification. We find an average VOT is 10.2 when a station has no POI access, and it is 7.1 when it has 3.1 POIs within 200m, or at the 10 percentile of the POI distribution. In column (3) of Table D1, we present a specification in which we include all the interactions in the price coefficient (our main specification only includes vehicle MSRP and the weekend dummy). The specification produces the average price coefficient of -0.15 with a standard deviation of 0.02, and the average duration coefficient of -0.99 with a standard deviation of 0.59. The average implied VOT is ¥6.9, or 28% of the average wage, with a median of 5.4 and a standard deviation of 4.6. Overall, these results from the alternative specifications show that our estimates from the main specification is robust to ad-hoc manipulation in modeling assumptions.

	Station Choice 10 Stations (1)	POI within 200m (2)	All Interactions (3)
Linear Parameters			
Pre-Charge Battery Volume (%)	-0.133 (0.018)	-0.115 (0.017)	-0.140 (0.018)
Log Vehicle MSRP (¥10,000)	0.171 (0.703)	-0.112 (0.988)	-0.016 (0.716)
Vehicle Battery Size (kWh)	0.094 (0.015)	0.080 (0.016)	0.105 (0.017)
Log Station Size	0.590 (0.048)	0.579 (0.049)	0.877 (0.083)
Log Nearby POI	0.120 (0.036)	0.092 (0.037)	0.002 (0.067)
Price sensitivity: $-\exp(\beta_0^p + \beta^p X_{ist}^p + \sigma^p \nu_i)$			
Constant	-1.532 (0.188)	-1.596 (0.212)	-1.631 (0.211)
Log Vehicle MSRP	-0.122 (0.061)	-0.139 (0.069)	-0.165 (0.101)
Log Nearby POIs	-	-	-0.025 (0.011)
Weekend Dummy	0.020 (0.039)	0.021 (0.044)	0.015 (0.040)
Log Station Size	-	-	0.058 (0.013)
Vehicle Battery Size (kWh)	-	-	0.002 (0.002)
σ^p	0.124 (0.018)	0.130 (0.020)	0.111 (0.018)
Time sensitivity: $-\exp(\beta_0^d + \beta^d X_{ist}^d)$			
Constant	-3.680 (1.687)	-3.547 (1.424)	-3.151 (1.440)
Log Vehicle MSRP	0.870 (0.701)	0.780 (0.583)	0.768 (0.610)
Log Nearby POIs	-0.139 (0.046)	-0.115 (0.039)	-0.097 (0.044)
Weekend Dummy	0.024 (0.547)	0.043 (0.199)	0.062 (0.189)
Log Station Size	-0.031 (0.055)	-0.026 (0.048)	-0.109 (0.051)
Vehicle Battery Size (kWh)	0.021 (0.013)	0.023 (0.012)	0.020 (0.013)
Dynamic Consideration			
End-of-Period Battery Volume (%)	0.121 (0.037)	0.133 (0.039)	0.091 (0.038)
End-of-Period Battery Volume^2/100 (%)	-0.096 (0.042)	-0.102 (0.041)	-0.059 (0.042)
End-of-Period Battery Volume^3/10000 (%)	0.034 (0.031)	0.042 (0.029)	0.016 (0.029)
Fixed Effects			
Day of Week FE	✓	✓	✓
Hour of Day FE	✓	✓	✓
Region (District) FE	✓	✓	✓
Price Instruments	✓	✓	✓
Implied VOT as Hourly Wage			
Mean	26%	28%	28%
Median	20%	21%	22%
Sample Size	18,785	18,785	18,785

Table D1: Utility Parameters

Note: Standard errors are reported in parentheses. MSRP, manufacturer's suggested retail price; VOT, value of time; FE, fixed effects; POI, point of interest.

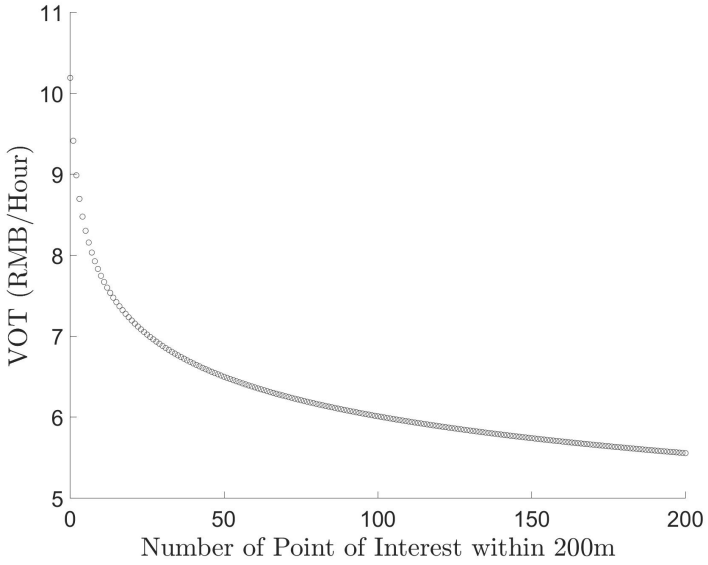


Figure D1: VOT Change in POIs within 200m

Note: The prediction is averages across drivers, stations and simulation draws for the random coefficient.

E Model Fit

In Figure E1, we present our model fit on the probability of making a station visit. The horizontal and vertical axes correspond to the pre-charge battery volume (in %) and the probability of charging, respectively. Uncolored dots represent data patterns and colored dots represent model predictions. To obtain the model predictions, we simulate 10,000 sequences of charging decisions during the sample period per driver, as detailed in Section 6. The model captures the decisions to visit stations at different battery conditions well.[51] In Appendix Figure E2, we show our model fit on consumer decisions over prices and charge duration. We slightly overestimate the probability that a driver chooses a lowest-price station (at 59.3% compared to 53.4% in the data), but the rest of the distributions look identical between the data and simulations.

[51]The small increase in predicted station visits near the battery volume of 100% arises because our continuation-payoff estimate is noisy in that range (because continuation payoff varies little across options because of the small charge volume). Because trips that end with a battery volume of 97% and above are rare (2% of the data), this prediction error will not affect our counterfactual outcomes.

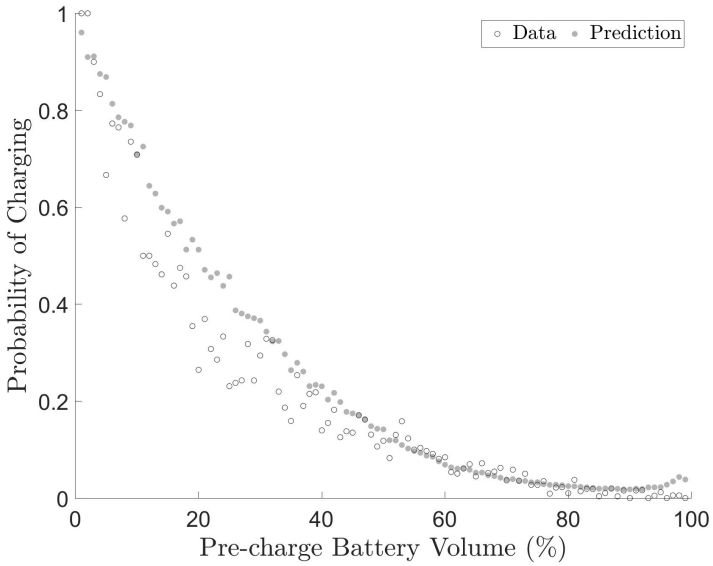
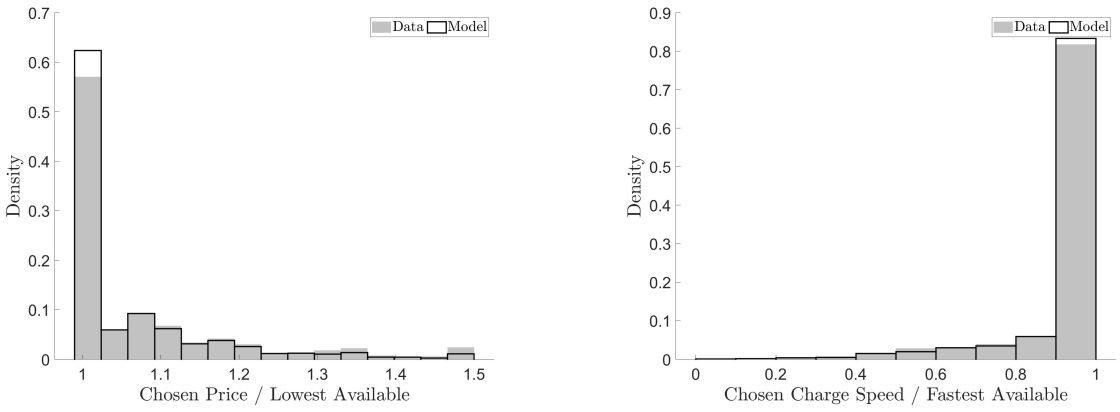


Figure E1: Data and Predicted Probability of Charging

Note: The empirical probability of charging corresponds to a frequency estimator at each pre-charge battery volume. The predicted probability of charging is based on 10,000 simulated sequences of charging events per driver, averaged across trips and drivers.



(a) Price Relative to Lowest

(b) Speed Relative to Fastest

Figure E2: Model Fit on Chosen vs Best in the Choice Set

Note: The histogram presents the price and the speed of the chosen station, relative to the best available in the choice set, conditional on a visit.

F Simulation Procedure

In Sections 6.1 through 6.4, we use simulated sequences of drivers' charging decisions to evaluate time costs and welfare. In this section, we provide details of these simulation analyses.

For the analyses in Sections 6.1 through 6.3, we compute the choice probability and draw action realization by drawing from $U(0,1)$. Throughout the simulation exercise, we assume $\xi_{ist} = 0$. The exact sequence of simulation is as follows.

1. Take driver i's first trip. Fix her random coefficient at the posterior value, i.e., $\mathbb{E}[\nu_i|d_{i1}, ...d_{iT_i}]$. Predict q_{ist} and the associated d_{ist} using the estimates from Tables 2 and 3. All the inputs to the utility are now observable besides ε_{ist}. Evaluate δ_{ist}, i.e., the utility from each option, excluding ε_{ist}.

2. If this simulation is for a counterfactual environment (i.e., either the number of available stations is limited, or some station characteristics are manipulated), then also compute δ_{ist} from the factual case (without manipulation). Compute $\max\{\delta_{ist}(\text{Counterfactual})\} - \max\{\delta_{ist}(\text{Factual})\}$, and add the number to the counterfactual outside option value.

3. Compute the choice probability, $Pr_{it}(s|q_{ist}, X_t, p_t)$.

4. Draw a random realization from $U(0,1)$, compare with the choice probability and determine the optimal action.

5. Compute the end-of-period charge volume. If the optimal action is to charge, add the charged volume to the current vehicle battery volume. Otherwise, add nothing. Proceed to the following trip. The remaining charge volume at the end of the following trip is given by "previous end-of-period charge volume - consumption during the trip".

The second step adjusts the level of outside utility in the counterfactual as relative to the maximum available inside utility. For instance, if a counterfactual reduces the utility from charging by 5 utils at the highest-utility station in the choice set, we also reduce the outside utility by 5. This adjustment thus fixes the probability of station visits between the factual and counterfactuals. This adjustment is necessary, because our reduced-form (polynomial) continuation payoff does not adjust automatically as we manipulate the inside-good utility: for instance, slower charging speed reduces the utility from inside goods, and thus, the frequency of station visits simply goes down unless we also adjust the outside utility. In the extreme scenario the driver may visit so few stations that she

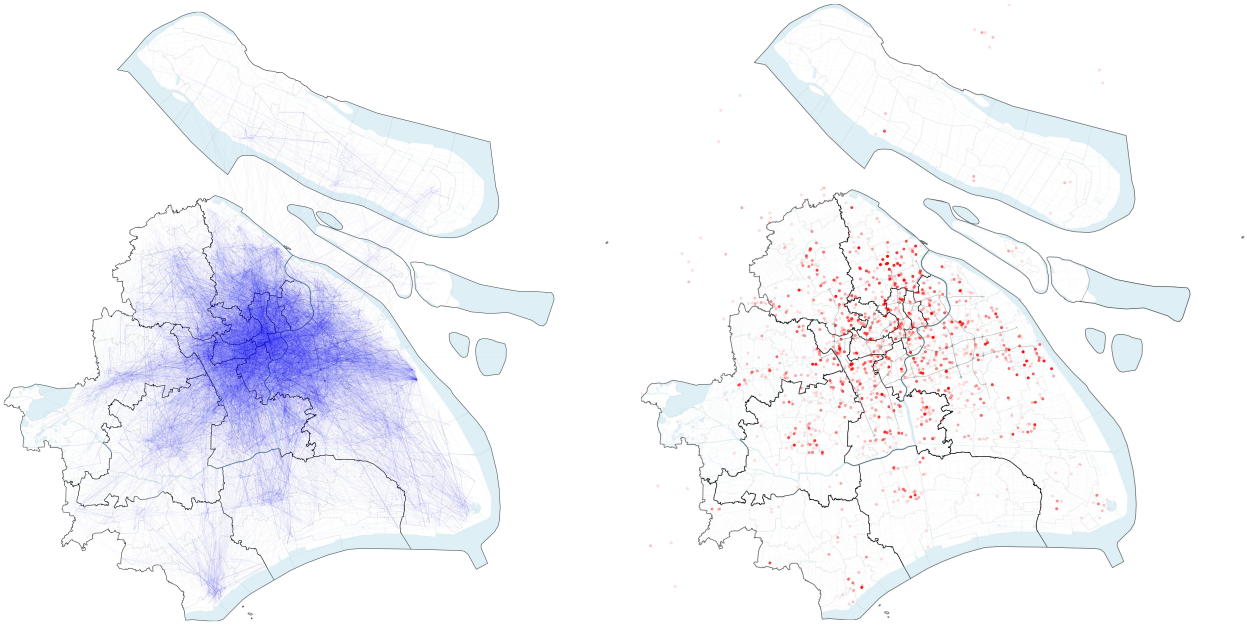
ends up with zero charge volume. To avoid this case, we also assume that the driver must charge the vehicle if her current battery volume cannot cover the subsequent trip, by assigning a large negative utility to outside option.

Analysis in Section 6.4 requires evaluation of driver's expected surplus, which includes realization of ε_{ist}. Thus, instead of drawing a realization of actions, we draw realizations of ε_{ist}, as follows:

1. Same as above.
2. Same as above.
3. Draw random realizations of ε_{ist} from $T1EV(0.5772, 1)$ for each option and add to δ_{ist}. The optimal action is the one that gives the highest utility.
4. Record $u_{is^*t} = \delta_{is^*t} + \varepsilon_{is^*t}$ of the chosen alternative. If this simulation exercise is a counterfactual and the outside option is selected, the recorded utility is $u_{is^*t} = \varepsilon_{is^*t}$, and not the utility adjusted in Step 2.
5. Same as above.

The reason we keep the zero outside utility in the record (and not the adjusted utility) is because this adjustment is a parsimonious way to account for lower *continuation payoffs*, and does not reflect the actual change in the utility from the outside option. This simulation exercise is computationally more intensive (because we draw multi-dimensional δ_{ist}'s rather than a single dimensional action), thus is only used in Section 6.4.

G Additional Figures and Tables



(a) Trip: origin to destination

(b) Fast charging locations

Figure G1: Description of Trips and Charging Locations

Note: Panel (a) presents trips (origin and destination connected by a straight line). We exclude trips that extend beyond the municipal boundary of Shanghai for visualization purposes. Panel (b) shows the locations of charging stations that drivers used during the sample period. Each dot denotes a visited station. Data source: the raw EV usage data with 40,382 trips and 6,293 charges.

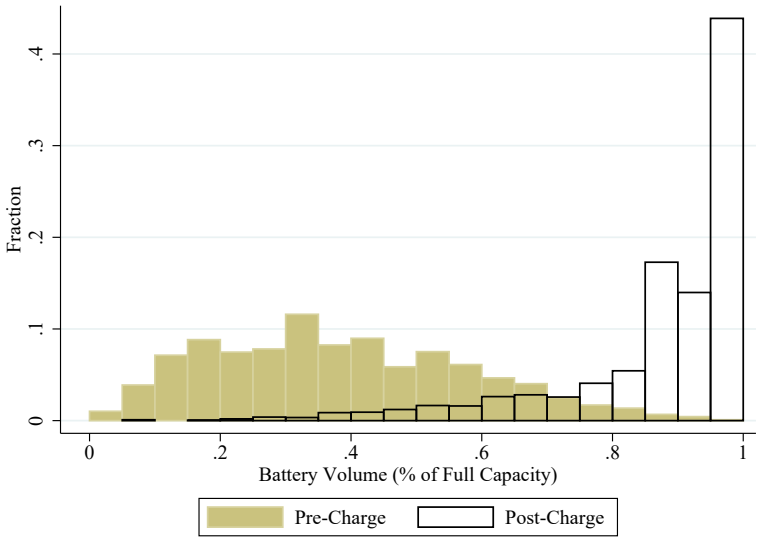
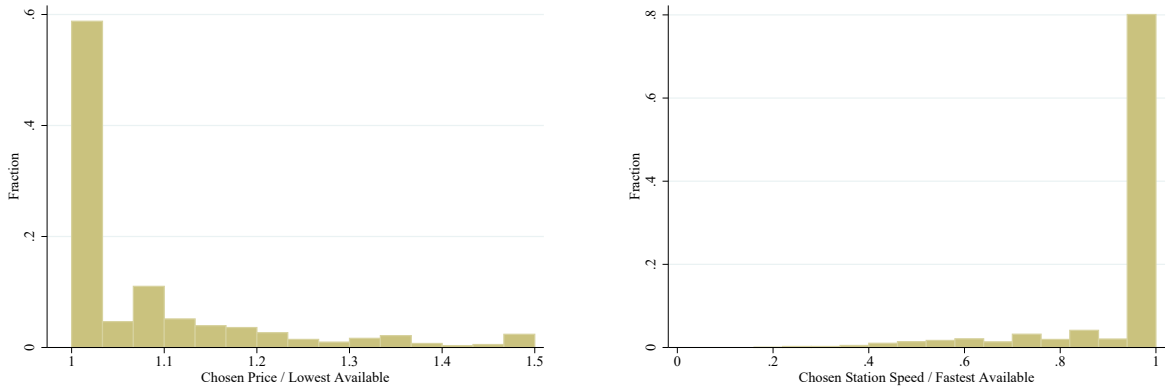


Figure G2: Pre- and Post-Charge Volume

Note: The unit of observation is a charge event. Yellow and white bars correspond to the pre- and post-charge battery volume, in % of vehicle battery size, respectively.

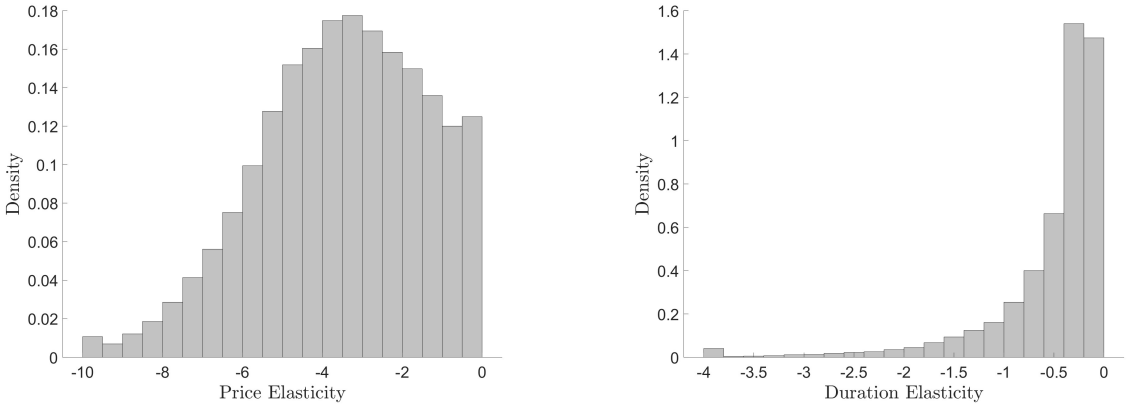


(a) Price Relative to Lowest

(b) Duration Relative to Fastest

Figure G3: Chosen vs Best in the Choice Set

Note: The histogram presents the price and the speed (available to the driver) of the chosen station, relative to the best available in the choice set, conditional on a visit.



(a) Own Price Elasticity (b) Own Duration Elasticity

Figure G4: Elasticity of demand with respect to Prices and Duration

Note: The unit of observation is a station-trip pair.

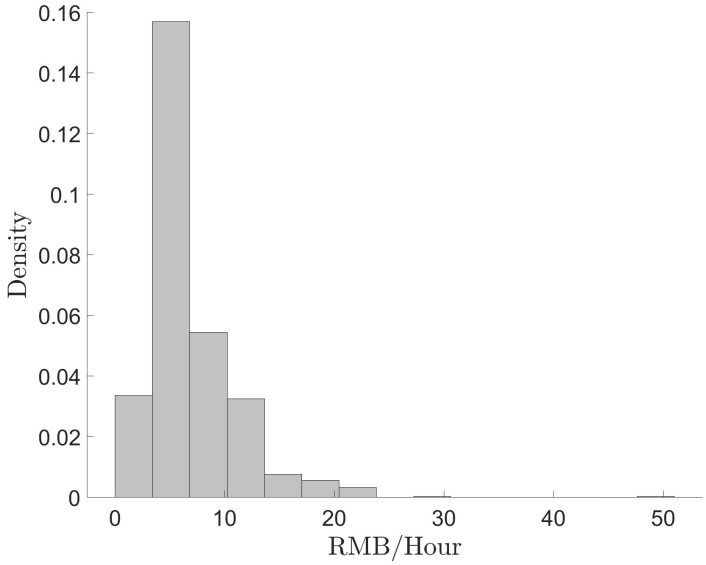


Figure G5: WTP for One Hour Reduction in Charge Time (Sample)

Note: The WTP is given by the ratio between the conditional mean of the duration and price coefficients. The unit of observation is a driver. The figure excludes one individual whose WTP is ¥150/hour.

	-25	-20	-15	-10	-5	Factual	5	10	15	20	25
Dec '23	144.66	164.45	179.46	192.27	203.29	212.68	217.96	222.41	225.64	228.30	230.39
Nov	132.74	152.45	167.57	180.27	191.20	200.45	205.75	210.24	213.50	216.22	218.37
Oct	131.87	150.03	164.13	176.16	186.51	195.52	200.56	204.83	207.93	210.51	212.54
Sep	127.29	144.73	158.36	169.92	179.85	188.51	193.33	197.43	200.39	202.87	204.83
Aug	109.99	129.51	143.69	155.41	165.32	173.84	178.53	182.53	185.41	187.79	189.68
Jul	87.39	107.22	121.55	133.38	143.28	151.82	156.58	160.62	163.49	165.87	167.74
Jun	64.16	85.09	100.26	112.78	123.11	132.11	138.70	143.40	146.59	149.17	151.17
May	37.81	58.41	72.92	84.82	94.65	103.23	109.65	114.21	117.28	119.77	121.68
Apr	22.40	41.97	55.67	66.84	75.94	83.90	90.02	94.26	97.19	99.58	101.40
Mar	1.09	27.32	42.77	54.45	63.65	71.54	77.75	82.04	85.04	87.51	89.34
Feb	-17.17	5.69	19.25	29.52	37.57	44.50	50.26	54.14	56.88	59.09	60.73
Jan '23	-19.19	3.31	16.60	26.62	34.48	41.20	46.81	50.60	53.25	55.39	57.00
Dec '22	-56.98	-35.08	-22.51	-13.35	-6.16	0.00	5.32	8.88	11.40	13.46	14.99
Nov	-80.84	-59.19	-46.92	-38.17	-31.33	-25.52	-20.31	-16.87	-14.49	-12.61	-11.20
Oct	-91.92	-70.95	-59.24	-50.97	-44.53	-39.13	-34.10	-30.82	-28.51	-26.72	-25.38
Sep	-101.47	-83.19	-73.46	-66.66	-61.23	-56.68	-52.22	-49.35	-47.38	-45.78	-44.53
Aug	-106.72	-89.52	-80.73	-74.83	-70.24	-66.51	-62.67	-60.37	-58.76	-57.42	-56.37
Jul	-118.72	-102.51	-94.59	-89.54	-85.69	-82.62	-79.10	-76.97	-75.48	-74.25	-73.29
Jun	-140.99	-125.23	-117.70	-112.91	-109.33	-106.49	-103.08	-101.06	-99.65	-98.53	-97.65
May	-148.23	-133.09	-126.06	-121.69	-118.45	-115.90	-112.69	-110.83	-109.57	-108.55	-107.77
Apr	-148.23	-133.09	-126.06	-121.69	-118.45	-115.90	-112.69	-110.83	-109.57	-108.55	-107.77
Mar	-148.64	-133.51	-126.48	-122.12	-118.89	-116.34	-113.13	-111.27	-110.02	-109.00	-108.22
Feb	-161.46	-146.95	-140.29	-136.23	-133.30	-130.98	-127.91	-126.21	-125.06	-124.13	-123.39
Jan '22	-163.13	-153.12	-148.60	-145.65	-143.44	-141.64	-138.90	-137.45	-136.49	-135.72	-135.12

Table G1: Welfare Table for Heatmap (Average, Figure 9)

	-25	-20	-15	-10	-5	Factual	5	10	15	20	25
Dec '23	195.90	213.88	227.29	238.18	247.09	254.50	259.31	263.14	265.04	266.51	267.49
Nov	179.46	196.89	210.24	221.00	229.82	237.12	241.84	245.66	247.55	249.07	250.09
Oct	176.52	192.43	204.91	215.17	223.49	230.49	234.97	238.59	240.29	241.68	242.60
Sep	167.83	182.88	194.73	204.42	212.29	218.92	223.15	226.63	228.29	229.61	230.51
Aug	154.13	170.45	182.57	192.33	200.21	206.78	210.87	214.21	215.80	217.04	217.91
Jul	121.12	137.52	149.60	159.34	167.08	173.54	177.58	180.96	182.59	183.83	184.66
Jun	92.08	108.95	121.39	131.41	139.27	145.83	149.89	153.34	154.98	156.25	157.12
May	59.18	75.92	88.11	98.00	105.68	112.05	115.99	119.31	120.94	122.18	123.02
Apr	39.65	54.95	66.06	75.06	82.05	87.81	91.41	94.44	96.00	97.19	97.98
Mar	8.32	32.14	45.75	55.73	63.18	69.19	73.03	76.22	78.05	79.40	80.35
Feb	-3.99	16.74	28.42	36.91	43.13	48.16	51.35	53.95	55.48	56.59	57.34
Jan '23	-5.16	15.38	26.85	35.11	41.15	46.01	49.14	51.68	53.14	54.20	54.91
Dec '22	-48.09	-28.41	-17.77	-10.16	-4.54	0.00	2.96	5.40	6.78	7.79	8.50
Nov	-77.73	-58.67	-48.52	-41.35	-36.06	-31.81	-28.99	-26.73	-25.39	-24.41	-23.74
Oct	-90.51	-72.33	-62.82	-56.11	-51.16	-47.18	-44.48	-42.35	-41.04	-40.08	-39.43
Sep	-103.87	-87.08	-78.68	-72.88	-68.54	-65.09	-62.79	-60.99	-59.93	-59.08	-58.50
Aug	-110.33	-94.56	-86.97	-81.88	-78.14	-75.24	-73.42	-72.02	-71.17	-70.52	-70.03
Jul	-127.02	-112.11	-105.33	-100.92	-97.74	-95.32	-93.73	-92.49	-91.71	-91.07	-90.60
Jun	-160.77	-145.95	-139.29	-134.97	-131.86	-129.49	-127.95	-126.74	-125.94	-125.33	-124.85
May	-166.17	-152.24	-146.27	-142.48	-139.79	-137.77	-136.42	-135.38	-134.70	-134.15	-133.72
Apr	-166.17	-152.24	-146.27	-142.48	-139.79	-137.77	-136.42	-135.38	-134.70	-134.15	-133.72
Mar	-166.58	-152.65	-146.67	-142.89	-140.20	-138.17	-136.82	-135.77	-135.10	-134.54	-134.12
Feb	-180.50	-167.49	-162.04	-158.68	-156.32	-154.54	-153.39	-152.50	-151.92	-151.43	-151.05
Jan '22	-185.12	-173.51	-168.84	-165.94	-163.91	-162.39	-161.39	-160.61	-160.09	-159.64	-159.30

Table G2: Welfare Table for Heatmap (MSRP Bottom Quartile, Figure 10 Panel (a))

	-25	-20	-15	-10	-5	Factual	5	10	15	20	25
Dec '23	104.64	126.05	142.15	156.12	168.77	180.05	185.31	189.96	193.69	196.64	198.83
Nov	99.16	120.87	137.17	150.95	163.54	174.64	179.96	184.69	188.40	191.44	193.69
Oct	97.99	118.60	134.03	147.34	159.49	170.81	176.20	180.87	184.61	187.70	189.96
Sep	94.34	114.48	129.92	143.02	154.95	166.12	171.39	175.97	179.64	182.69	184.94
Aug	66.12	91.61	108.88	122.96	135.30	146.46	151.63	156.07	159.64	162.47	164.61
Jul	49.93	76.57	94.41	108.99	121.74	133.23	138.67	143.18	146.65	149.52	151.64
Jun	30.55	59.68	79.79	96.33	110.58	123.22	129.61	134.94	138.82	142.00	144.35
May	20.27	48.58	67.62	83.26	96.84	109.05	115.36	120.61	124.35	127.45	129.80
Apr	4.47	32.07	50.77	65.91	78.74	90.33	96.53	101.47	105.02	108.08	110.38
Mar	-26.08	14.88	37.39	53.79	66.91	78.45	84.93	90.10	93.80	97.02	99.37
Feb	-38.09	-2.70	17.12	31.59	43.22	53.44	59.45	64.21	67.67	70.59	72.72
Jan '23	-39.86	-5.11	14.34	28.52	39.85	49.76	55.59	60.29	63.66	66.52	68.61
Dec '22	-85.72	-50.76	-31.82	-19.03	-8.88	0.00	5.35	9.69	12.83	15.52	17.41
Nov	-103.49	-68.78	-49.74	-37.01	-26.96	-18.31	-12.85	-8.31	-5.17	-2.65	-0.87
Oct	-114.88	-80.56	-61.94	-49.55	-39.94	-31.84	-26.60	-22.26	-19.18	-16.77	-15.05
Sep	-121.32	-91.75	-76.79	-67.03	-59.15	-52.51	-48.08	-44.25	-41.67	-39.54	-37.93
Aug	-118.65	-90.47	-76.77	-68.36	-61.88	-56.63	-53.35	-50.61	-48.59	-46.86	-45.52
Jul	-130.77	-103.14	-89.92	-82.27	-76.53	-71.98	-69.06	-66.51	-64.65	-63.06	-61.85
Jun	-143.42	-116.77	-104.25	-97.08	-91.83	-87.69	-84.96	-82.55	-80.79	-79.38	-78.28
May	-149.89	-123.83	-111.76	-104.98	-100.05	-96.18	-93.67	-91.43	-89.83	-88.50	-87.49
Apr	-149.89	-123.83	-111.76	-104.98	-100.05	-96.18	-93.67	-91.43	-89.83	-88.50	-87.49
Mar	-150.37	-124.31	-112.24	-105.46	-100.53	-96.66	-94.15	-91.91	-90.32	-88.99	-87.97
Feb	-159.28	-133.42	-121.56	-115.03	-110.44	-106.86	-104.60	-102.58	-101.18	-100.01	-99.06
Jan '22	-161.02	-143.56	-136.00	-131.43	-128.13	-125.44	-123.69	-122.08	-120.97	-120.05	-119.29

Table G3: Welfare Table for Heatmap (MSRP Top Quartile, Figure 10 Panel (b))