The Trade-Offs of Curbside Parking: Evidence from Demand-Based Pricing*

Hoang-Anh Nguyen† Tuyetanh L. Tran‡

Job Market Paper

November 5, 2025

[Please click here for the latest version]

Abstract

City governments face a trade-off in managing curb space: providing parking to facilitate access to consumption amenities and generate revenue, versus allocating it to alternative land uses. In this paper, we quantify the welfare implications of curbside parking and evaluate alternative policies for managing curb space through parking instruments. We develop a structural model of drivers' joint destination and parking decisions: drivers choose which destination to visit under imperfect information about parking availability, then decide where to park near the chosen destination. We estimate the model using high-frequency data on metered parking transactions and GPS data on visits to points of interest in San Francisco, one of the few cities that have implemented demand-based pricing for curbside parking. We find that, while drivers value curbside parking, the present discounted value of parking revenue and driver surplus generally falls short of local assessed land values, which proxy for the economic value of land uses. Compared to a revenue-maximizing uniform pricing scheme, San Francisco's demand-based pricing generates about 30% more revenue while reducing cruising trips by nearly 70%. Our counterfactuals show that reducing parking supply by roughly 6% and lowering the status quo demand-based prices by \$1.25 citywide preserves parking welfare, with only a modest revenue loss, while freeing curb space for other uses.

*Authors are listed alphabetically. We are deeply grateful to Katja Seim, Steven Berry, and Cody Cook for their guidance, kindness, and continuous support throughout the project. We also received valuable feedback from Phil Haile, Charles Hodgson, Victoria Marone, Robert Miller, Jingyuan Wang, Kevin Williams, and many seminar participants. We acknowledge and appreciate staff at Dewey and the San Francisco Municipal Transportation Agency (SFMTA) for providing data support. We thank Alexiy Sukhenko, Kenya Wheeler, and Hank Wilson at SFMTA and Mary Catherine Snyder at Seattle Department of Transportation for institutional insights. We also thank Tuan Phung-Duc for helpful discussions on queueing theory and Songyuan Teng for computational advice.

†Tepper School of Business, Carnegie Mellon University; Email: hoangann@andrew.cmu.edu.

‡Department of Economics, Yale University; Email: tuyetanh.tran@yale.edu.

1 Introduction

Urban transportation plays a critical role in supporting economic activity, especially through improving traffic externalities and expanding market access to destinations (Redding and Turner, 2015). While recent empirical work has made progress in quantifying the economic impacts of transportation systems, [1] it has been focused on the movement of people and goods and paid less attention to parking infrastructure, which complements these moving transportation systems. In the United States, curbside parking, parking that occupies street curb space, accounts for a substantial share of the parking infrastructure. For example, in New York City, street parking occupies nearly 14 square miles, about the area of 13 Central Parks (Grabar, 2024). On the one hand, curbside parking helps expand market access to urban amenities through convenient vehicle storage and generates public revenue. On the other hand, mismanagement of curbside parking can lead to inefficient use of urban land. When underpriced, curbside parking also contributes to traffic externalities such as congestion and emissions, as many drivers have to *cruise* to look for parking. These traffic externalities can intensify residents' opposition to new housing and commercial constructions (Kashner and Ross, 2025), highlighting the broader implications of parking policies for urban development.

In this paper, we quantify the welfare effects of curbside parking, measured by the sum of city revenue and driver surplus minus time costs of cruising, and evaluate alternative policies for managing curb space through parking instruments. Our analysis focuses on metered parking in San Francisco, the first U.S. city to adopt citywide demand-based parking pricing.[2] Demand-based pricing sets meter rates by time and location in response to past demand, mitigating excess demand in some parking blocks. We begin by describing the empirical setting and presenting summary statistics on the composition of urban amenities (e.g., restaurants, shops) in San Francisco, travel patterns in the city, and its parking conditions. Motivated by the descriptives, we develop and estimate a structural model of drivers' joint decisions of where to go and where to park, incorporating how parking considerations factor into travel decisions. San Francisco's demand-based parking pricing creates a feedback mechanism: prices affect parking availability, availability influences destination choices, and these choices shape parking demand, which in turn feeds back into prices. Our model disentangles these interactions to estimate drivers' preferences for parking and the extent to which parking considerations affect their travel choices. Using the framework, we quantify drivers' value of consumption amenities and parking, and provide a measure of the wel-

[1] For example, Durrmeyer and Martínez (2024), Cook and Li (2025), and Cook et al. (2025) study how congestion pricing and tolling policies mitigate traffic externalities; Tsivanidis (2023) examines how public transit improvements expand commuter market access and increase social surplus; Almagro et al. (2024) evaluates alternative urban transportation policies and their welfare and distributional implications.

[2] Demand-based parking pricing is considered one of the most significant reforms in parking policies (Pierce et al., 2015).

fare of curbside parking. Finally, we compare the effects of alternative curbside parking policies, such as price adjustments and parking supply reductions, to evaluate their trade-offs.

We combine multiple data sources to link drivers' destinations with parking choices: GPS mobile-device visit data from Veraset, and San Francisco administrative data, including parking inventory records, high-frequency meter transactions, travel surveys, and records on street characteristics and restrictions. From the GPS data, we observe drivers' home Census block group, points of interest (POIs) that they visit, and the timestamps of those visits.[3] From the administrative data, we obtain information on the availability and hourly prices of each parking block near a destination, as well as other street characteristics that influence parking preferences.[4] Although we do not directly observe who parks at each block, we recover drivers' relevant parking choice sets using the geolocation of destinations and parking blocks. We augment these geolocations with assumptions about the radius around a destination where drivers look for parking, consistent with the urban planning literature on parking (Millard-Ball et al., 2020; Weinberger et al., 2020). Our primary sample covers the universe of San Francisco metered parking transactions in June 2019, as well as visits to San Francisco by Bay Area drivers. In our analysis, we focus on visits and parking activities in the city's commercial areas, a region containing roughly 60% of San Francisco's POIs and metered parking spaces despite occupying only 20% of the city's land area.

We begin by showing key patterns in travel and parking in San Francisco. About half of all visits to consumption amenities in the city by Bay Area residents are made by private car. Unlike commuters who often have reserved workplace parking, visitors to these destinations rely on curbside spaces or commercial lots. In San Francisco's commercial areas, where curb spaces are usually restricted and off-street parking is costly and sparse, fewer than 2% of establishments provide on-site parking, making metered parking particularly salient for drivers to these areas. We document rich variation in meter rates and availability across times of day and adjacent locations, with blocks charging more than \$6/hour located next to blocks priced at \$1–2, suggesting strong heterogeneity of parking choices. On average, most blocks have high parking availability thanks to demand-based pricing. Nonetheless, at peak parking times (12 pm to 3 pm), up to 5% of all metered blocks in the city have almost no available parking spots during a three-hour period. The observed variation in parking conditions allows us to estimate drivers' elasticities with respect to different parking attributes, which is generally difficult in settings with uniform parking prices. Finally, we find that the present discounted value (PDV) of parking revenue captures only a small fraction of local assessed land values, which proxy for the economic value of alternative land uses. On average, the PDV of a block's parking revenue is about \$400,000, and the ratio of this PDV to the underlying land value is roughly 20%. This evidence suggests that current metered parking

[3] A POI is a well-defined establishment, such as a restaurant or a supermarket.
[4] A parking block is the street segment between a pair of opposing blockfaces (see Figure A.1).

generates modest monetary returns relative to the value of the land it occupies. This comparison, however, reflects only the revenue side of parking's economic return, omitting the surplus drivers derive from using these spaces.

To quantify the value drivers derive from curbside parking within their travel decisions and to evaluate how they would respond to counterfactual policies, we develop a structural model of drivers' joint destination and parking choices. In the first stage, each driver chooses a destination within San Francisco's commercial areas to visit. This decision depends on destination characteristics and expectations about parking outcomes once they arrive. In the second stage, drivers choose where to park based on parking locations' characteristics (e.g., hourly prices, proximity to destination, steep terrain) and under imperfect information about parking availability.[5] Upon arrival at the preferred metered parking block, if it has no open spot, the driver is assumed to resort to parking at an outside option (i.e., off-street garage). This redirection of parking due to metered parking unavailability is interpreted as *cruising* in our framework. The two-stage travel decision is represented by a nested logit model. We close the model by specifying how drivers form rational expectations about parking availability at each metered block. We build on the Erlang B queuing formula (Erlang, 1909) to link the likelihood of finding an open parking space to expected parking demand and supply: as more drivers attempt to park at a block with a fixed number of spaces, the block's availability declines.

We follow Train (2009) and Azar et al. (2022) and estimate the nested logit model sequentially. We first treat destination choices as fixed and estimate drivers' preferences for parking to obtain each destination's inclusive value of parking. We then estimate destination choices given these values and quantify the role of parking in drivers' travel decisions. We address two sources of endogeneity. In the lower-level parking choice, meter rates and expected availability may be correlated with unobserved parking quality, such as safety and convenience. We instrument for prices and availability using a combination of BLP-style instruments based on nearby parking characteristics and data on temporary non-metered parking closures (e.g., for street cleaning). In the upper-level destination choice, the value of nearby parking may be correlated with unobserved destination preferences, as popular destinations tend to have more crowded parking. We address this by using moderately distant parking disruptions from special events, such as festivals or parades, in the city. When a special event takes place near but not immediately adjacent to a destination, it does not directly impact foot traffic to that destination. However, the event can create temporary local parking congestion, and if the distance between the event and the destination's relevant parking locations is moderate, this adverse parking condition can spill over to parking near the destination, reducing its availability and, in turn, influencing drivers' destination choices through parking considerations.

[5] This is consistent with parking fees and availability predictions publicly available on the website of the San Francisco Municipal Transportation Agency and through parking apps, such as PayByPhone, SpotAngels, and Parknav.

We find that drivers are responsive to both prices and expected availability of parking, suggesting that pricing is an effective tool in managing parking demand. At the destination level, drivers are most sensitive to proximity to their home, but their consideration of expected parking outcomes ranks second, more salient than preferences for other destination attributes such as the composition of establishments. This is consistent with parking being a subsequent yet significant component of the travel decision. From these estimated preferences, we quantify driver surplus for each metered parking block. Building on the earlier comparison of the PDV of meter revenue to underlying land value, we now incorporate driver surplus into the analysis. When accounting for driver surplus, we find that while most blocks' parking surplus, the sum of revenue and driver surplus, remains below land value, about 5.5% exceed it, with some exceeding by a large margin. On average, the PDV of parking surplus amounts to roughly 40% of underlying land value. This suggests substantial spatial heterogeneity in curbside parking surplus and that drivers can derive considerable surplus from metered parking. These findings underscore the importance of designing flexible, location-specific parking policies, consistent with the mechanism underlying demand-based pricing.

Finally, we conduct two sets of counterfactual exercises to quantify the role of demand-based pricing in managing curbside parking and to evaluate alternative policy designs. First, we compare the status quo demand-based pricing with a revenue-maximizing uniform pricing scheme. The status quo generates nearly 30% more revenue than the uniform scheme while reducing cruising trips by roughly 69%, saving drivers more than $500,000 in time costs each month. Second, we find that San Francisco could reduce its parking supply by roughly 6% and still achieve the same total welfare, defined as parking surplus minus time costs of cruising, by lowering prices by $1.25/hour citywide. This price adjustment compensates drivers for the decline in availability due to reduced supply, while causing only modest changes in revenue and cruising costs. Together, these findings highlight the importance of coordinating curb space provision and parking prices in managing and achieving more flexible use of curb space.

Related Literature. Our paper integrates ideas from several strands of literature in economics, urban planning, and operations research. A central area of focus is the economics of parking, which studies how parking policies affect traffic externalities and urban welfare. Theoretical work on parking dates back to Vickrey (1954), who proposed spatially and temporally differentiated parking prices. Foundational contributions by Arnott and Rowse (1999), Anderson and De Palma (2004), Arnott and Inci (2006, 2010), and Arnott et al. (2015) provide theoretical support for parking pricing as a tool to reduce traffic externalities and improve urban welfare. Despite the well-established theoretical foundation, empirical evidence on parking pricing remains limited. Recent empirical work using event-study designs finds evidence that parking pricing policies can reduce congestion and related traffic externalities from private cars.[6] Krishnamurthy and Ngo (2020) analyzes data

[6]See Kashner and Ross (2025) for evidence of traffic externalities from car usage.

from San Francisco's pilot demand-based pricing program and concludes that the implementation of demand-based pricing substantially increased bus ridership and reduced traffic flow, yielding significant economic gains through lower congestion and emissions. Ostermeijer et al. (2022) studies a citywide increase in street parking prices in Amsterdam, a policy that raised prices by 66% on average and made the city the most expensive in the world for curbside parking, and finds that the policy reduced street parking demand by nearly 20% and led to a noticeable reduction in traffic flow. Gragera et al. (2021) uses a reduced-form framework that incorporates both observed parking fees and unobserved cruising costs across proximate parking locations to estimate the price elasticity of parking demand, informing our BLP-style structural approach that captures substitution among nearby parking options. Research in urban planning and operations research complements this economic literature by emphasizing the design, operation, and optimization of parking systems.[7] We contribute to this interdisciplinary literature on parking by quantifying the welfare of curbside parking within a joint destination-parking choice framework and evaluating the trade-offs of alternative policy designs for managing curb space through parking instruments.

Our research is also tied to the literature on spatial equilibrium in transportation. Recent empirical work, such as Frechette et al. (2019), Rosaia (2020), Buchholz (2022), Almagro et al. (2024), Barwick et al. (2024), Durrmeyer and Martínez (2024), Castillo (2025), Cook and Li (2025), and Cook et al. (2025), has focused on in-motion vehicle systems, including ride-sharing platforms, traffic flows, public transit networks, and tolls and road pricing. However, parked vehicles are equally critical, with an average car spending 95% of its lifetime parked (Shoup, 2018). We contribute to the literature by studying parking decisions and quantifying the equilibrium and welfare implications of vehicle storage, offering a complementary perspective on the still state of vehicles and their role in facilitating urban mobility and expanding market access to urban destinations.

Our paper also contributes to the growing literature that uses GPS data to study urban mobility and spatial behavior. Athey et al. (2018), Cao et al. (2024), Cook (2024), and Couture et al. (2025) use GPS-based data to estimate heterogeneous preferences for amenities and compute travel time; Gupta et al. (2022) and Cook and Li (2025) analyze the welfare and distributional implications of new transportation systems; Almagro et al. (2024) studies optimal urban transportation policies, emphasizing the trade-offs between road pricing and public transit policies; Miyauchi et al. (2025) examines how consumption externalities shape urban agglomeration. We extend this literature by combining GPS-based amenity visit data with high-frequency parking transactions to study the

[7] Some key studies include Shoup (2005, 2018), Kelly and Clinch (2009), Millard-Ball et al. (2014); Millard-Ball et al. (2020), Chaniotakis and Pel (2015), Lehner and Peer (2019), Dalla Chiara and Goodchild (2020), Weinberger et al. (2020), and Feldman et al. (2022). Feldman et al. (2022) also studies the welfare effects of demand-based pricing, but our paper differs in several important ways: i) we allow parking considerations to influence destination choices; ii) we observe granular destination visits and can align them with high-frequency parking data; and iii) we use data from San Francisco's citywide implementation of demand-based pricing rather than data from the pilot program, where demand-based pricing was introduced only in limited areas.

joint destination–parking travel decision and quantify the extent to which parking considerations factor into travel decisions.

Finally, our work also builds on the literature on amenity choice, which examines the welfare distribution and spatial sorting of urban amenities. Some notable contributions include Couture (2016), Davis et al. (2019), Su (2022), Cook (2024), and Almagro and Domínguez-Iino (2025). We extend this literature by building on Cook (2024) to introduce a second stage to the amenity choice framework, in which drivers decide where to park after selecting their destinations, thereby capturing the interaction between destination preferences and nearby parking availability.

2 Pricing Curbside Parking in the U.S.

Since the introduction of parking meters in the early twentieth century, pricing has become a core instrument of curbside parking policy, complementing rationing such as time limits, permits, special zoning, and other regulatory tools. For decades, most cities relied on simple pricing schemes that set flat, and often low, hourly rates across time and space. These traditional pricing schemes have been criticized for encouraging over-parking, misallocating scarce curb space away from drivers with high willingness to pay (WTP), and inducing cruising. In contrast, variable and market-based pricing strategies have gained growing support (Anderson and De Palma, 2004; Arnott and Inci, 2006; Shoup, 2005, 2018). These schemes adjust rates by time and location in response to observed past demand or increase prices with the duration of parking, which better balances parking demand and supply (Kaufman et al., 2012). In this paper, we focus on demand-based pricing schemes that regularly update meter rates by time and location, a reform regarded as one of the most significant in modern parking (Pierce et al., 2015).

2.1 Demand-Based Pricing

Demand-based pricing, also known as dynamic or performance-based pricing, refers to a pricing mechanism that adjusts prices periodically to reflect demand. In the context of parking, the conceptual foundation traces back to Vickrey (1954), which proposes replacing time limits and fixed, low curbside prices with a system of higher rates, comparable to off-street parking, that vary across time and space. Under this system, parking spaces are effectively "rented" according to economic principles, balancing demand with scarce supply, preventing overparking, and reducing cruising. Demand-based pricing helps ensure that curb spaces remain available for drivers with the highest WTP rather than being allocated on a first-come, first-served basis.

Shoup (2005) proposes a practical rule of thumb for setting demand-based parking prices: adjust prices to maintain a persistent one or two open spaces per parking block. In practice, local governments implement this by keeping the average curbside utilization within a preset target range. Each period, usually monthly or quarterly, local governments measure the occupancy of metered parking. If a parking segment's occupancy falls within the target utilization range, prices remain unchanged; otherwise, they are adjusted accordingly. The specific pricing rules and occupancy thresholds vary across local jurisdictions.

2.2 Demand-Based Pricing in San Francisco

Among the U.S. cities using demand-based pricing for parking, San Francisco offers one of the most comprehensive implementations.[8] Launched as a pilot program between 2011 and 2013 under the name SFpark Pilot, it was funded through an $18 million federal grant from the U.S. Department of Transportation. The program's primary goal, aligned with Shoup (2005), was to use pricing to maintain some open spots per parking block, thereby reducing cruising for spaces, a significant contributor to urban congestion and greenhouse gas emissions. To achieve this goal, every parking block's occupancy rate was targeted to be between 60% and 80% through periodic price adjustments. Besides reducing traffic externalities, by stabilizing parking conditions, the city sought to improve market access to local establishments, thereby boosting foot traffic and sales for urban amenities. The pilot operated in seven parking management districts, including Civic Center, Downtown, Fillmore, Fisherman's Wharf, Marina, Mission, and South Embarcadero, collectively covering roughly 6,000 metered spaces, about 25% of the city's total on-street parking (SFMTA, 2014b). Evaluations reported improved parking availability and reduced cruising, congestion, and greenhouse gas emissions in treated areas. Following the pilot success, San Francisco implemented a citywide demand-based pricing system in early 2018, becoming the first U.S. city to adopt such an approach at scale.

Before the program, San Francisco, like many other U.S. cities, relied on flat meter rates. These rates remained constant throughout days and years, with downtown areas charged $3.50/hour and other districts charged $2.00/hour (Krishnamurthy and Ngo, 2020). Under the demand-based pricing program, prices increase (decrease) quarterly in $0.25 increments when occupancy rates exceed (fall below) the target range, with the price cap as of June 2019 set at $8.00/hour and the minimum rate at $0.50/hour. These price adjustments are announced publicly on the San Francisco Municipal Transportation Agency (SFMTA) website, and in general, price information is available on parking apps, such as PayByPhone. Pricing varies by parking block, day type, and time of day. A parking block is the street segment between a pair of opposing blockfaces (see Figure A.1). Day types include weekdays (Monday through Friday) and weekends (Saturday), while meters are free

[8]Other cities implement demand-based pricing at a more limited scale. For example, Seattle uses zone-level rather than block-level pricing, Los Angeles implements demand-based pricing only in selected neighborhoods, and Boulder varies prices by location but not by time.

on Sundays. A typical day is divided into several time bands: morning (before 12 pm), early afternoon (12 pm to 3 pm), late afternoon (3 pm to 6 pm), and, in some areas, evening (after 6 pm). As detailed by Pierce and Shoup (2013), this pricing scheme steers short-stay and time-sensitive drivers (including those with mobility limitations or dependents) toward more convenient spaces, while nudging others to park slightly farther away.

3 Data

To construct our analysis data, we rely on several raw datasets. This section describes our primary data sources and provides the summary statistics for our main variables.

3.1 Data Sources

3.1.1 Parking Data

We combine San Francisco's administrative data on parking inventory, street characteristics, and metered parking transactions to construct our main dataset for parking.

The parking inventory data, provided by SFMTA, includes information on the geographic coordinates of each parking location. For metered parking, the dataset reports the number of metered spaces per parking block, each meter's hourly rate at different times, and other applicable restrictions. We focus on metered spaces that are open for the general public parking.[9] We aggregate the data from meter-level to parking-block-level as meters on the same parking block generally share the same prices and restrictions. For off-street parking, the data generally do not include information on prices or supply but indicate whether each facility is open to the public or reserved for private use. In this paper, we refer to facilities open to the general public as public-use garages (or off-street parking), and those restricted to specific users as private-use garages.[10]

We augment the parking inventory data with information on each street segment's characteristics that impact parking convenience and desirability, such as whether the segment is one-way or has steep terrain.[11] A street segment's orientation (one-way or two-way) affects parkers' ease of access to the segment, while a segment with steep slope may be less attractive to parkers due to more difficult parking maneuvers and walking. Data on street orientation and steep terrain are obtained from the San Francisco Department of Public Works (SFDPW) and Department of Planning

[9] Some metered spaces are reserved for specific uses (e.g., loading zones, accessible parking), and some spaces are operated under parking permits.

[10] Public-use garages may be operated by either the local authority or private entities. They include off-street parking facilities such as surface lots and multi-story parking structures that are accessible to the general public, rather than limited to specific groups such as employees, tenants, or customers.

[11] San Francisco has many hilly streets.

(SFDP), respectively.[12] We also obtain each street segment's sweeping schedule from SFDPW and later exploit its exogeneity as an instrument when estimating parking preferences.[13] Since street cleaning in the commercial areas of San Francisco occurs only during the early part of the day, we estimate the model using data from timebands 1 (opening to 12 pm) and 2 (12 pm to 3 pm). We hereafter refer to timebands 1 and 2 as morning and afternoon parking, respectively.[14]

We then match the data with SFMTA's high-frequency metered parking transactions from June 1 to June 30, 2019. The transaction data records the specific meter at which each transaction occurs, along with its timestamp and duration. From this information, we infer each parking block's availability, that is, the fraction of time during which open spaces exist within a timeband. We then combine the number of transactions at each parking block with information from other datasets to derive each block's market share. See Section 6.2 for more details.

3.1.2 Destination Data

Our destination data comes primarily from Veraset mobile-device visit data covering the period from June 1 to June 30, 2019.[15] Veraset collects raw GPS coordinates and timestamps from anonymized devices, and then aggregates these sequences into visits to POIs. Veraset visit data reports each visit's POI geolocation and category (based on NAICS code classifications), chain affiliation (brand), on-site parking information, timestamp, minimum dwell time at the POI, anonymized device ID, and the device's inferred home census block group (CBG).[16] We focus on visits to restaurants, shops, and entertainment venues, as these destinations tend to generate metered parking demand. We follow Cook (2024) in choosing relevant NAICS codes and subcategories within the main categories.[17] We aggregate the data from POI-level to street-block-level to avoid the issue of POIs with zero visits within a given time unit (i.e., a specific timeband on a particular day). We augment the visit data with information on special events occurring near each street block in San Francisco during June 2019, obtained from SFMTA.

We focus on visits by Bay Area residents, including those living in San Francisco, to destinations in San Francisco, as this sample includes drivers who look for short-term parking and find

[12] SFDP provides geographic data on steep terrain, defined under the California Environmental Quality Act as slopes exceeding 20%.

[13] See Section 6.2.

[14] Average daily visits to our consumption amenities of interest in the commercial areas are approximately 9,000, 10,000, and 11,500 for timebands 1, 2, and 3, respectively. Thus, restricting estimation to the first two timebands still provides a representative sample of parking behaviors.

[15] We access the Veraset visit data (Veraset, 2022) through Dewey (deweydata.io), a data access platform for academic researchers.

[16] Some POIs lack reported information on on-site parking. In these cases, we supplement the POI data with SFMTA off-street parking data to fill missing values.

[17] See Table A.1 in Cook (2024) for more details.

metered parking a relevant option.[18] While the GPS-based visit data allows us to infer the total number of visits to POIs in the city, we do not observe which visits are made by private cars. To recover the proportion of visits by car, we augment the visit data with SFMTA 2019 Travel Decision Survey. The survey was conducted among Bay Area residents who visit San Francisco. Respondents were asked about their trip purposes, modes of transportation, and broadly defined home locations. From this survey, we derive the share of Bay Area residents who drive to San Francisco destinations. We infer short-term parkers from the survey based on respondents' reported trip purposes. For example, respondents who drive to San Francisco for shopping or dining are likely short-term parkers, whereas those who drive to work or school are likely not. Combining this with the visit data, we can estimate the population of drivers and parkers in the city.

3.1.3 Assessed Land Value Data

Our land value data comes from the San Francisco Office of the Assessor-Recorder (SF Assessor-Recorder). The dataset reports assessed land values and parcel sizes for each fiscal year, as recorded for property tax purposes. We use data from the fiscal year 2018-2019 in our analysis. We compute the assessed land value per square foot for each parcel and aggregate to the neighborhood level to obtain the average assessed land value per square foot per neighborhood.[19] Assessed land values serve as a proxy for the economic value of alternative land uses and provide a standardized measure of land value that is comparable across locations in the city.

3.2 Summary Statistics

We present summary statistics for the key variables used in the analysis.

Figure 1 shows the map of San Francisco with POIs that include restaurants, shops, and entertainment venues. We focus on the city's commercial areas (inner city), shown in blue, located along the northeastern corridor and the waterfront, areas with the highest density of economic activity. These areas occupy only about 20% of San Francisco's land area but account for roughly 60% of POIs in our sample. The inner city consists of the following neighborhoods: Financial District/South Beach, South of Market, Mission, Castro/Upper Market, Nob Hill, Western Addition, Marina, Japantown, Chinatown, North Beach, Russian Hill, Pacific Heights, and Tenderloin.

Table 1 presents the characteristics of destinations in the inner city. Our sample includes 8,814 POIs aggregated into 1,057 street blocks. Restaurants constitute the largest category, representing nearly half of all POIs in the sample, while shops and entertainment venues also account for substantial shares. Notably, only less than 2% of POIs in the sample have their own parking lots.

[18] Short-term parkers refer to drivers who look for hourly parking, rather than daily, monthly, or permit parking.

[19] We use the Analysis Neighborhood definition provided by the San Francisco Mayor's Office of Housing and Community Development (MOHCD) to define a neighborhood.

Figure 1: Map of San Francisco

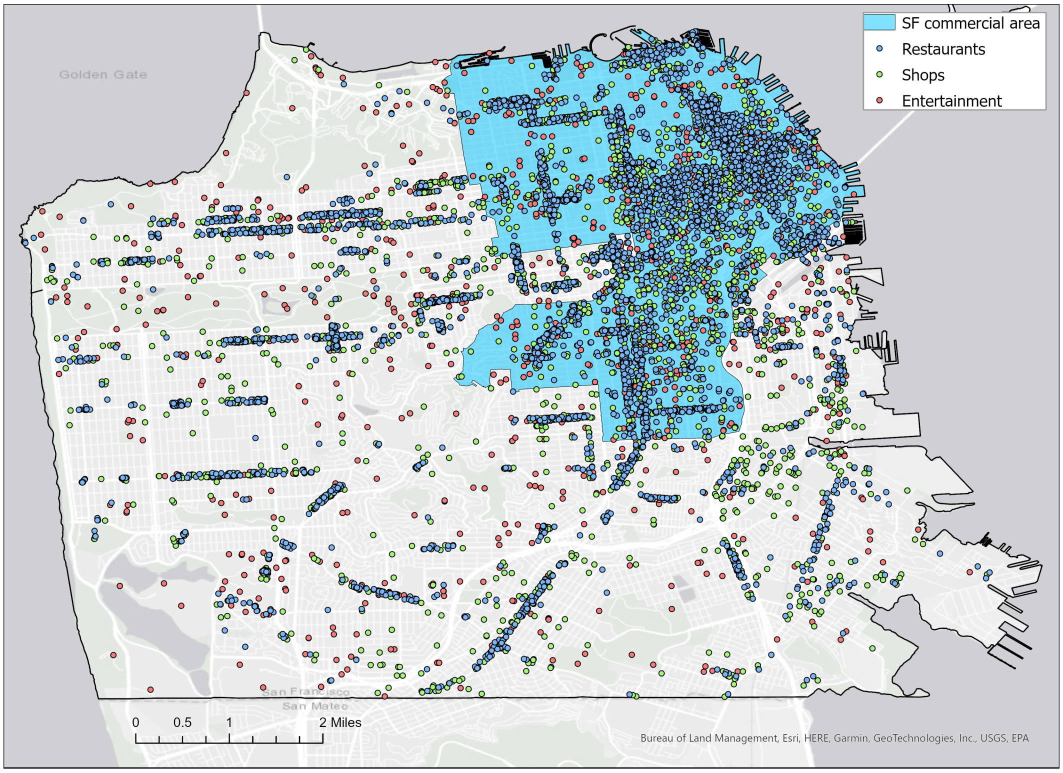

Notes: The map shows the distribution of POIs in San Francisco, color-coded by amenity category: restaurants (blue), shops (green), and entertainment venues (orange). The NAICS codes for these POIs are provided in Table A.1 in Cook (2024). The shaded blue area represents the city's commercial neighborhoods, located along the northeastern corridor and waterfront. The POI data is sourced from Veraset for the period June 1 - 30, 2019.

Figure 2 shows price variations for the 886 metered blocks overseen by SFMTA in the inner city included in our sample. Panel (a) shows average hourly prices across blocks, averaged over morning and afternoon periods as well as weekdays and weekends. There is substantial cross-sectional variation: hourly prices range from \$0.50 to over \$7.00 in June 2019, with most blocks priced between \$2 and \$4. Panel (b) shows the difference between morning and afternoon prices across blocks. A considerable number of blocks have afternoon prices that are substantially higher than morning prices (up to over \$5). This rich variation in prices, both across location and over time, allows us to estimate price elasticity in our later analysis, which is generally difficult in settings with uniform pricing.

Table 2 shows the summary statistics for other key parking variables, such as parking availability and temporary non-metered parking closures (e.g., street sweeping), at the block–timeband–date level. Thanks to demand-based pricing, parking availability is generally high: on average, a block has an open space 96% of the time, meaning the block has no available spaces for only about 10

Table 1: POI Characteristics

$N = 8,814$	Percentage
Amenity Categories	
Restaurants	47.3%
Shops	36.7%
Entertainment	16.1%
Parking Lot	
Own Parking Lot	1.7%
No Parking Lot	98.3%

Notes: This table reports characteristics of POIs in San Francisco's commercial areas, focusing on restaurants, shops, and entertainment venues. The POI data is sourced from Veraset (June 2019). We show the distribution of POIs by amenity category and the share of POIs with on-site parking lots. The sample includes 8,814 POIs across 1,057 street blocks. Information on POI parking lots is obtained from Veraset and supplemented with SFMTA data on off-street parking for private use to fill in missing values.

Figure 2: Hourly Price Variation across Location and Time

(a) Cross-Sectional (b) Temporal

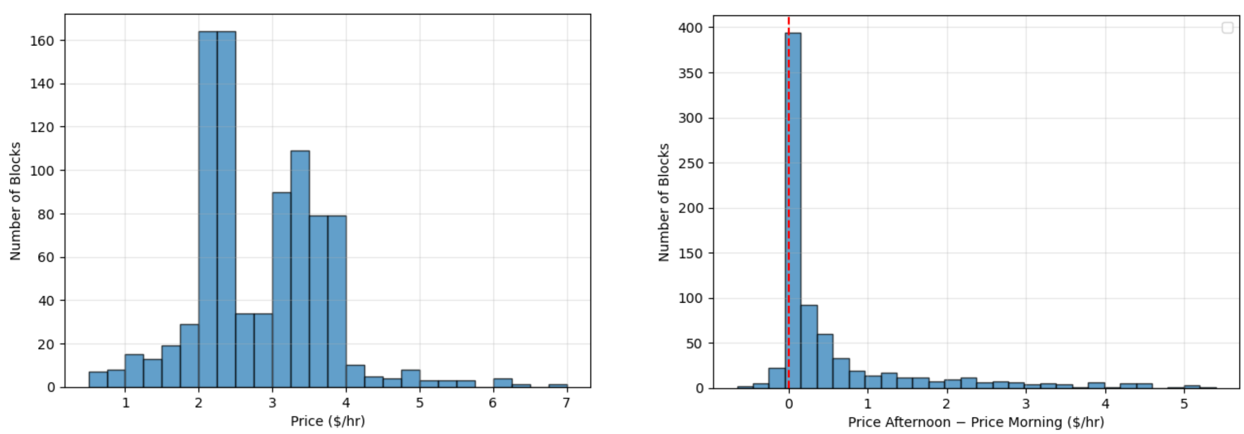

Notes: The figure presents the distribution of hourly parking prices across the 886 metered blocks overseen by SFMTA in San Francisco's commercial areas for June 2019. Panel (a) shows the cross-sectional distribution of average hourly prices across block, where each block's price is averaged over morning and afternoon periods as well as weekdays and weekends. Panel (b) shows the difference between afternoon and morning prices across block.

minutes per 180-minute timeband. However, availability varies substantially across blocks, timebands, and dates, with some blocks occasionally experiencing near-zero availability. Non-metered parking closures also display significant variation: while the average block experiences roughly four closures per timeband, some blocks have none and others experience as many as 83 closures. Together, the statistics underscore the rich variation in parking conditions across San Francisco's commercial areas.

Table 2: Summary Statistics of Parking Availability and Non-Metered Parking Closures

Variable	Mean	Std. Dev.	Min	Max
Parking Availability (%)	96.16%	13.50%	0.50%	100%
Non-Metered Parking Closures (occurences)	4.26	10.73	0	83

Notes: This table presents summary statistics for parking availability and non-metered parking closures at the block–timeband–date level in San Francisco's commercial areas for June 2019. Parking availability is measured as the fraction of minutes within a 180-minute timeband during which at least one space is open. Non-metered parking closures include non-metered parking blocks closed for street cleaning. Statistics are calculated across 41,210 block-timeband-date observations.

4 Descriptive Statistics

In this section, we present some descriptive statistics about the empirical setting and discuss how they motivate our structural model in Section 5.

Travel Patterns and Demand for Metered Parking. The prevalence of driving to consumption amenities is particularly relevant for metered parking demand. Table 3 summarizes travel patterns to San Francisco by Bay Area residents based on the 2019 Travel Decision Survey. Leisure trips, including dining out, entertainment, recreation, shopping, and errands, account for nearly 40% of all trips to the city. Among these, roughly half are made by private car, either driving alone or with others. Unlike commuters who often have reserved workplace parking, visitors to restaurants, shops, and entertainment venues rely on parking options such as street parking, commercial garages, or POI-owned lots. As shown in Table 1, only fewer than 2% of POIs provide on-site parking. Moreover, in the commerical areas, most curb spaces are subject to tow-away and other restrictions, and off-street garages are costly and sparse, making metered parking vital in accommodating visitors to consumption amenities and facilitating access to San Francisco's commercial neighborhoods.

Spatial Heterogenenity in Parking Conditions. In the inner city, parking conditions vary substantially even across adjacent blocks. For example, blocks charging more than $6/hour can be located next to blocks priced less than $2, which indicates strong heterogeneity in parking choices and suggests that drivers make highly granular parking decisions. Figure 3 illustrates the spatial variation in hourly parking prices and availability across blocks, zooming in on the North Beach neighborhood. Parking availability and prices differ significantly among nearby blocks. Availability tends to be lower where prices are high, consistent with the demand-based pricing mechanism. However, some blocks with relatively high prices also exhibit high availability despite being surrounded by POIs. This heterogeneity suggests that drivers may face trade-offs when choosing

Table 3: Travel Patterns

	Percentage
Trip Purpose	
Leisure Trips	39.6%
Commuting Trips	57.8%
Other	2.6%
Trip Mode	
Private Car	48.4%
Other Modes	51.6%

Notes: This table presents the distribution of trips made by Bay Area residents to San Francisco, based on the 2019 Travel Decision Survey. We report the share of leisure trips (dining out, entertainment, recreation, shopping, and errands) versus commuting trips (work-, school-, and home-related travel), as well as the share of private car trips (where respondents drove alone or with others) versus other transportation modes (e.g., ride-hailing (Uber/Lyft), taxi, public transit, bicycle, walking).

where to park, taking into account factors such as proximity to destinations, parking prices, or physical attributes of each block. This observation underscores the need for a model that captures drivers' choices and disentangles their preferences at a granular level, as significant variation in parking outcomes can arise even across neighboring blocks.

Comparison of Parking Revenue and Land Value. We then compare metered blocks' parking revenues with their underlying land values to assess the relative economic return of parking use. For each parking block, its land value is estimated by multiplying the neighborhood land value per square foot by the standard metered stall size and the number of metered spaces on the block. We calculate the PDV of annual parking revenue, using a 7% discount rate, to make the two measures comparable.[20]

Figure 4 plots the PDV of parking revenue against land value for all metered blocks in the inner city. Land values vary substantially across space, and parking revenue is positively correlated with land value, as expected. However, all parking blocks fall below the 45-degree line, indicating that parking revenue typically captures only a small fraction of the land's potential economic value. On average, each block generates an annual meter revenue of about $30,000, which translates into a PDV of more than $400,000 and accounts for roughly 23% of the underlying land value. This comparison, however, reflects only the local government's gain of parking's economic return, omitting the surplus drivers derive from using these spaces. To account for driver surplus and to quantify a more comprehensive parking return, we present a structural model in the next section.

[20]The 7% discount rate follows Circular A-94: Guidelines and Discount Rates for Benefit-Cost Analysis of Federal Programs.

Figure 3: Heatmap of Metered Parking

Notes: The figure presents a spatial heatmap of parking availability and hourly prices for a section of the North Beach neighborhood. Each colored strip represents a metered parking block. The top panel shows parking availability by block, and the bottom panel displays hourly parking prices. Blue dots represent the locations of POIs (restaurants, shops, and entertainment venues). The parking data is obtained from SFMTA for June 2019, and the POI locations are sourced from Veraset for the same period.

5 Model

5.1 Model of Driver Choice

We develop an economic model of drivers' joint decisions on destinations and parking locations, incorporating how parking considerations factor into travel decisions. In the first stage, Bay Area drivers draw their trip purposes (dining, shopping, or entertainment), then they choose a street block in San Francisco that has amenities in that category (hereafter, destination). Street blocks in San Francisco outside the inner city are treated as a single outside option. In the second stage, they decide where to park near the chosen destination in the inner city. Drivers have imperfect information about parking availability, if they arrive and no curbside spots are open, they resort to using a public-use garage. This redirection of parking due to metered unavailability is interpreted as cruising in our framework. This two-stage decision is represented by a two-level nested

Figure 4: Parking Revenue and Land Value

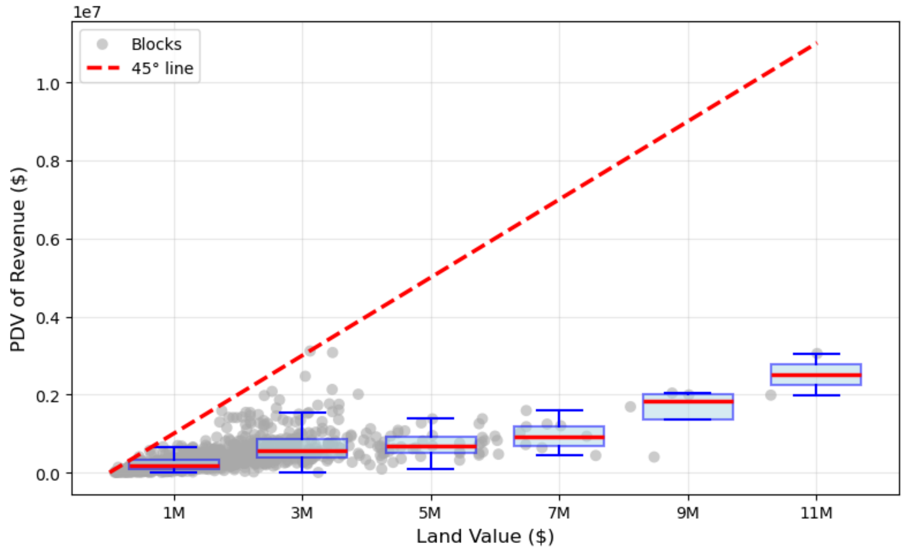

Notes: The figure presents the relationship between parking revenue and land value for all metered blocks in San Francisco's commercial areas. The horizontal axis shows the estimated land value for each parking block, calculated as the neighborhood land value per square foot multiplied by the standard metered stall size and the number of metered spaces on the block. The vertical axis shows the PDV of annual parking revenue, calculated using a 7% discount rate. Each gray dot represents a parking block, and the red dashed line denotes the 45-degree line, where parking revenue equals land value. The box plots show the median and interquartile range of parking revenue within each land value bin. Parking revenue is computed using SFMTA metered parking transactions for June 2019, and neighborhood land value per square foot is computed using data from SF Assessor-Recorder.

logit (Cardell, 1997; McFadden, 1978), where each destination forms a nest, and parking locations nearby are alternatives within nests. The nesting structure captures the feedback between destination and parking choices.

5.1.1 Stage 2: Parking Demand Conditional on Destination

Given drivers' destinations in the inner city, we model their parking choices. We define a market as timeband t on date d. We perform our analysis at the timeband level rather than at a finer resolution (e.g., the minute level) for two reasons. First, the price data is defined at the timeband and day type level: within a given day type, the three-hour morning and afternoon timebands have their own prices, which SFMTA sets based on aggregate demand for each timeband. Second, aggregating to the timeband level reduces computational burden. The set of parking locations accessible to a driver visiting destination j in market td is $\mathcal{K}_{jtd}$, which consists of all metered blocks and public-use garages within 400 meters of the destination. The radius used to define the parking choice set follows established research in urban planning, such as Millard-Ball et al. (2020) and Weinberger et al. (2020). All nearby public-use garages are included as a single outside option (denoted 0) in this stage.

Upon arriving at j, driver i has two parking options: i) park directly at the outside option, or ii) attempt to park at metered block $k \neq 0 \in \mathcal{K}_{jtd}$ first, and if k is unavailable, resort to the outside option. We assume that public-use garages are always available, and thus attempting to park at a garage is equivalent to directly parking at that garage.

A driver visiting j during timeband t on date d receives the following mean utility from parking at the outside option

$$\delta_{0|jtd} = \gamma_t + \alpha_{garage} PG_j \tag{1}$$

with γ_t being timeband fixed effects and PG_j being the number of public-use garages around j.

We define the mean utility of *parking* at metered block k as

$$\delta_{ktd} = \gamma_{n(k)} + \mathbf{X}_{ktd}^{park} \boldsymbol{\alpha}_{\boldsymbol{X}} \tag{2}$$

where $\gamma_{n(k)}$ is neighborhood fixed effects and $\mathbf{X}_{ktd}^{park}$ is a vector of parking block k's characteristics,[21] including hourly parking price P_{ktd} relevant for timeband t on date d, street orientation, and an indicator for steep terrain. We abuse notation by letting P_{ktd} denote prices at different locations and times, but it is important to note that prices vary by block, timeband, and day type (weekday and weekend) rather than by specific date.

When deciding whether to try parking at block k, a driver cares about the distance between their destination j and parking block k, the mean utility of parking at k that would be realized if k is available, and the mean utility that they would receive when parking at the outside option in case k is unavailable. Driver i thus receives the following mean utility from *attempting to park* at metered block k

$$\delta_{k|jtd} = \alpha_{dist} Dist_{kj} + \underbrace{EA_{ktd} \times \delta_{ktd}}_{k \text{ available}} + \underbrace{(1 - EA_{ktd}) \times \delta_{0|jtd}}_{k \text{ unavailable}} + \xi_{ktd} \tag{3}$$

where $Dist_{kj}$ is the straight line distance between parking block k and destination j normalized by the radius of parking choice set,[22] EA_{ktd} is the expected availability of k during timeband t on date d, which we define in Section 5.2, and ξ_{ktd} denotes unobserved preferences for attempting to park at block k.

The value of parking near destination j is the expected utility from choosing the best parking

[21] Drivers care about neighborhood characteristics in deciding where to park. For example, within 400 meters of a destination, drivers can access parking blocks in both Financial District/South Beach and Tenderloin neighborhoods, but they might prefer the former as the latter has a higher crime rate.

[22] The straight line distance is measured as the distance between the centroids of destination j and parking block k.

location to attempt, conditional on visiting j. In the nested logit literature, it is referred to as a nest's *inclusive value* and has a closed-form expression

$$V_{jtd}^{park} = \log\left[\exp\left(\delta_{0|jtd}/\beta_{park}\right) + \sum_{k\in\mathcal{K}_{jtd}/\{0\}} \exp\left(\delta_{k|jtd}/\beta_{park}\right)\right] \tag{4}$$

where β_{park} is a nesting parameter that links parking and destination choices (see Equation 6 below).

5.1.2 Stage 1: Destionation Demand Conditional on Trip Purpose

During timeband t on date d, driver i draws an amenity category c and decide which street block with amenity c to visit. The set of street blocks in San Francisco is denoted $\mathcal{J}$, where $0 \in \mathcal{J}$ represents the outside option consisting of destinations in San Francisco outside the commerical areas.

If driver i chooses the outside option, they receive the normalized utility

$$U_{i0ctd} = \tilde{\gamma}_{h(i)} + \tilde{\gamma}_c + \tilde{\gamma}_t + \varepsilon_{i0ctd} \tag{5}$$

where the $\tilde{\gamma}$ terms denote home CBG ($h(i)$), amenity category (c), and timeband (t) fixed effects.

Driver i cares about the distance between their home and destination j, the value of parking near j, and other attributes of j. The utility they receive when visiting destination $j \in \mathcal{J}/\{0\}$ is

$$U_{ijctd} = \underbrace{\beta_{dist}Dist_{h(i)j} + \beta_{park}V_{jtd}^{park} + \beta_{brand}Brand_{jc} + \mathbf{X}_{jdt}^{dest}\boldsymbol{\beta}_{\mathbf{X}} + \omega_{jctd}}_{\delta_{j|h(i)ctd}} + \varepsilon_{ijctd} \tag{6}$$

where $Dist_{h(i)j}$ is the straight line distance between driver i's home and destination j,[23] V_{jtd}^{park} is the parking value around j defined in Equation 4, $Brand_{jc}$ is the number of branded POIs of category c on block j, $\mathbf{X}_{jdt}^{dest}$ contains other attributes of j (including time-varying factors like nearby special events within a certain radius of j and time-invariant characteristics like store composition),[24] and ω_{jctd} denotes unobserved preferences for j conditional on category c.[25] The idiosyncratic error terms ε follow appropriate distributions to generate the nested logit structure as in Cardell (1997).

[23]The straight line distance is measured as the distance between the centroid of a destination and the weighted centroid of the home CBG. Because home locations are observed only at the CBG level, which is coarser than the destination (street block) level, we use a weighted centroid to provide a more consistent measure of distance.

[24]In our main analysis, we define a nearby special event as one occurring within 400 meters of a destination, but our results are robust to a wide range of radii.

[25]Controlling for store composition (i.e., the number of restaurants, shops, and entertainment venues on block j) is a reduced-form way to account for trip-chaining behaviors, similar to how Cao et al. (2024) models consumer utility as a function of density.

We let $\delta_{j|h(i)ctd}$ denote the mean utility of choosing destination j conditional on home CBG $h(i)$ and drawn trip purpose c.

5.1.3 Choice Probabilities

Stage 1. From Equation 6, conditional on coming from home CBG $h(i)$ and drawing amenity category c, the probability that driver i visits destination j takes the standard logit formulation (McFadden, 1973)

$$s_{j|h(i)ctd} = \frac{\exp(\delta_{j|h(i)ctd})}{\sum_{j'} \exp(\delta_{j'|h(i)ctd})} \tag{7}$$

From the conditional choice probability $s_{j|h(i)ctd}$, we compute the share of drivers visiting destination j during timeband t on date d, denoted s_{jtd}. See Appendix B.1 for details on the computation.

Stage 2. From Equation 3, the probability that a driver arriving at $j \in \mathcal{J}/\{0\}$ attempts parking at block $k \in \mathcal{K}_{jtd}$ is

$$s_{k|jtd} = \frac{\exp(\delta_{k|jtd}/\beta_{park})}{\sum_{k' \in \mathcal{K}_{jtd}} \exp(\delta_{k'|jtd}/\beta_{park})} \tag{8}$$

The mean utility $\delta_{k|jtd}$ is normalized by the parameter β_{park} from Equation 6 to account for the nesting structure, which captures the correlation among parking alternatives within the same destination nest.

See Equation 14 or Appendix B.1 for how to compute s_{ktd}, the unconditional share of drivers attempting to park at metered block k upon visiting the inner city, from the conditional share $s_{k|jtd}$.[26]

5.2 Expected Availability of Parking

As noted in Section 3.1.1, we measure a block's availability during a 180-minute timeband as the fraction of minutes with at least one open space. This measure can be interpreted as the probability of finding an open space at a randomly chosen minute during that timeband. Before arriving at a parking block, drivers hold expectations about this probability, which, in practice, is formed with their experience of local parking conditions or via parking apps' predictions.

[26] Equation 14 shows how s_{ktd} is computed when destination shares are observed in the data, whereas Appendix B.1 derives s_{ktd} fully from the model, which we use to compute new equilibria in the counterfactual analysis in Section 8.

Specifically, the ex-post availability of block k within a 180-minute timeband is measured as

$$A_{ktd} = \frac{o_{ktd}}{180} \tag{9}$$

where o_{ktd} is the number of minutes within 180-minute timeband t that block k has at least one open parking spot. We observe o_{ktd}, and thus A_{ktd}, from the metered transaction data.

In our model, we assume o_{ktd} follows a binomial distribution

$$o_{ktd} \sim Binomial(180, EA_{ktd}) \tag{10}$$

where EA_{ktd} is the ex-ante probability that block k has at least one open spot in any given minute during timeband t on date d, representing drivers' expected availability. The expectation of availability is rational, since

$$\mathbb{E}[A_{ktd}] = \frac{\mathbb{E}[o_{ktd}]}{180} = \frac{180 \times EA_{ktd}}{180} = EA_{ktd} \tag{11}$$

We model the expected availability of parking block k during timeband t on date d using a binomial generalized linear model (GLM) with a logit link function

$$\text{logit}(EA_{ktd}) = \eta_A A_{kt} + \breve{\mathbf{X}}_{ktd}\boldsymbol{\eta}_{\boldsymbol{X}} + \breve{\gamma}_d \tag{12}$$

where $\breve{\gamma}_d$ denotes date fixed effects, A_{kt} is the predicted baseline availability of block k during timeband t, and $\breve{\mathbf{X}}_{ktd}$ represents observable shocks to parking block k during timeband t on date d that may impact availability. This specification captures how drivers form expectations in practice: they have baseline knowledge of how busy block k typically is during timeband t (captured by A_{kt}), which they update for each date based on observable factors (captured by $\breve{\mathbf{X}}_{ktd}$ and $\breve{\gamma}_d$).

We predict the baseline availability A_{kt} using the Erlang B formula (Erlang, 1909). In our parking context, Erlang B gives the probability that all parking spaces on block k are occupied during timeband t when a driver arrives, given the average rate at which drivers attempt parking at k (demand) and the number of parking spaces at block k (supply), both inferred from our parking data. The Erlang B formula is described in detail in Appendix B.2. A_{kt} is then one minus the computed Erlang B.

We include the number of nearby non-metered blocks that are closed for parking during timeband t on date d in $\breve{\mathbf{X}}_{ktd}$. Non-metered blocks are typically reserved for special uses, such as loading zones, accessibility parking, or permit parking. When these blocks are closed due to temporary restrictions, such as street cleaning, drivers who would have used these blocks might shift to nearby metered blocks, lowering availability there.

6 Estimation

The estimation of drivers' travel decisions with destination and parking choices in Section 5.1 is done sequentially following the approach in Train (2009) and Azar et al. (2022). We estimate the model in Section 5 in three steps. First, we estimate Equations 10 and 12 to obtain the expected availability parameters ($\boldsymbol{\eta}$) and predict $\widehat{EA}_{ktd}$. Second, we use $\widehat{EA}_{ktd}$ and the observed parking characteristics to estimate the parking preference coefficients $\boldsymbol{\alpha}$ and the nests' inclusive values $\hat{V}_{jtd}^{park}$ from Equation 4. Note that at this stage, we can only recover $\tilde{\boldsymbol{\alpha}} = \dfrac{\boldsymbol{\alpha}}{\beta_{park}}$, where β_{park} is the nesting parameter (see Section 6.2 for further details). Finally, from $\hat{V}_{jtd}^{park}$ and the observed destination attributes, we estimate the destination preference parameters ($\boldsymbol{\beta}$).

6.1 Estimation of Expected Availability

We estimate the parameters $\boldsymbol{\eta}$ and the fixed effects $\breve{\gamma}_d$ in Equation 12 using Maximum Likelihood Estimation. The corresponding log-likelihood is given by [27]

$$\log \mathcal{L} = \sum_{k,t,d} [o_{ktd} \log(EA_{ktd}) + (180 - o_{ktd}) \log(1 - EA_{ktd})] \tag{13}$$

where EA_{ktd}, as defined in Equation 12, is a function of $\boldsymbol{\eta}$, $\breve{\gamma}_d$, $\breve{\mathbf{X}}_{ktd}$, and A_{kt}. The variables $\breve{\mathbf{X}}_{ktd}$ and A_{kt} are either directly observed or inferred from the parking data.[28]

We report the full estimation results in Table A.2. We then convert the estimates to average marginal effects for interpretation. Overall, a one percentage point (pp) increase in baseline availability A_k corresponds to a 0.21 pp increase in expected availability EA_{kt}. This suggests that drivers' expectations are strongly anchored to baseline conditions, with roughly one-fifth of variation in baseline availability translating into changes in expected availability. Closures of nearby non-metered parking blocks have a modest yet statistically significant effect (at the 1% level) on expected availability, reducing expected availability by about 0.28 pp when at least one nearby non-metered parking block is closed. Date fixed effects account for the remaining variation.

6.2 Estimation of Parking Preferences

In estimating the parking choice model, we treat drivers' destination choices as given from the data. Specifically, we observe s_{jtd}^{inner}, the share of drivers visiting destination j relative to the total number of drivers visiting the inner city.

[27] Here we omit terms that do not depend on $\boldsymbol{\eta}$ in the log-likelihood.

[28] See Appendix B.2 for details on how to compute A_{kt} from the parking data.

From the model, we compute $s_{k|jtd}$ (see Equation 8), the probability that a driver arriving at destination j in the inner city attempts to park at metered block k. The total share of drivers visiting the inner city who attempt to park at metered block k, regardless of their specific destination, is then

$$s_{ktd} = \sum_{j \in \mathcal{J}/\{0\}} s_{k|jtd} \times s^{inner}_{jtd} \tag{14}$$

To estimate the parking parameters, we equate the model-predicted shares s_{ktd} with their observed counterparts in the data. We now discuss how to infer this variable from the parking and destination data described in Section 3.

First, we use GPS data and travel survey data to infer the total number of Bay Area drivers visiting San Francisco's commercial areas during each timeband on each day in June 2019. From SFMTA's high-frequency metered parking transactions and inventory data, we observe two key variables for each block k: i) the number of parking transactions, and ii) block-level availability for each timeband and date A_{ktd}, where availability is defined as the share of minutes with at least one open spot over the timeband length of 180 minutes (see Section 5.2). We then compute the share of drivers who successfully park at block k, sh^{park}_{ktd}, by dividing the number of transactions at k by the total number of Bay Area drivers visiting the inner city in market td from the destination data.

Nonetheless, we need to infer the share of drivers who *attempt to park* at block k, not just those who successfully find a spot. We recover this by adjusting for availability: the share of drivers attempting to park at k during timeband t equals the share successfully parking at k divided by the probability of finding an open spot there, i.e., the availability of k during timeband t.[29]

$$sh^{attempt}_{ktd} = \frac{sh^{park}_{ktd}}{A_{ktd}} \tag{15}$$

We then equate s_{ktd} and $sh^{attempt}_{ktd}$ to estimate the parking parameters. Note that in choosing where to park, drivers face different parking trade-offs conditional on their destinations. For example, they have different parking distances to k or different numbers of public-use garages to resort to as their outside option, depending on their chosen destination j. This destination-specific variation creates heterogeneity in parking utilities across drivers that is analogous to demographic interactions in random coefficients models. We thus estimate the normalized parking parameters $\tilde{\boldsymbol{\alpha}}$ following the nested fixed point algorithm in Berry et al. (1995) for random coefficient estimation. We implement the estimation with PyBLP, using the procedure in Conlon and Gortmaker (2020). Although we do not observe which driver parks where, we know the share of visits to destina-

[29] This approach to recovering attempt shares follows the discussion of parking utilization in Feldman et al. (2022).

tions s_{jtd}^{inner} and the share of parking attempts at each block s_{kdt}. We can treat destination-specific variables as drivers' demographic characteristics and use s_{jtd}^{inner} as the weight for simulated agents associated with destination j.

Identification. Expected availability EA_{ktd} and hourly prices P_{ktd} are likely correlated with the unobservables ξ_{ktd}. Desirable unobserved characteristics, such as safety and convenience, attract more drivers to block k, reducing expected availability and triggering higher demand-based prices. We address this endogeneity with instrumental variables (IVs) for expected availability and prices.

We instrument for prices using BLP-style IVs: exogenous characteristics of neighboring parking blocks. These characteristics affect block k's equilibrium price but do not directly enter drivers' utility for parking at k, satisfying the relevance and exclusion restriction conditions of valid IV (Berry et al., 1995). Specifically, we use neighboring parking blocks' street orientation, steep terrain indicators, and the characteristics of POIs near the neighboring blocks.

To instrument for expected availability, we exploit variation from two sources. First, the capacity (number of parking spaces) of block k affects expected availability through the baseline availability A_{kt} computed from the Erlang B formula (Section 5.2). Blocks with more spaces have higher baseline availability, all else equal. Second, nearby non-metered parking block closures (e.g., for street sweeping) increase EA_{ktd} by displacing parking from non-metered to metered blocks. These satisfy the relevance condition for valid instruments. Both instruments also satisfy the exclusion restriction because drivers condition their choices on expected availability itself, not directly on the underlying capacity nor the closure variables that influence this availability. Capacity provides cross-sectional variation, while non-metered parking closures add temporal variation in our instruments.

The set of IVs $\boldsymbol{Z}_{ktd}$ provides the orthogonality condition $\mathbb{E}[\boldsymbol{Z}_{ktd}\xi_{ktd}] = 0$, which is used in the Generalized Method of Moments estimation following the procedure of Conlon and Gortmaker (2020).

Parameter Estimates. Note that this estimation procedure gives us $\tilde{\boldsymbol{\alpha}} = \dfrac{\boldsymbol{\alpha}}{\beta_{park}}$, the parking parameters normalized by the nesting coefficient. Column (2) of Table 4 reports the estimates of $\tilde{\boldsymbol{\alpha}}$. The results generally align with our expectations. Parkers dislike the distance between their destinations and parking blocks, one-way street segments, and high expected prices. The coefficient on steep terrain is statistically insignificant, indicating that drivers generally do not mind parking on steeper segments. The value of the outside option (conditional on the chosen destination) increases with the number of public-use garages nearby.

We later estimate β_{park} and recover the true parking parameters $\boldsymbol{\alpha}$ in Section 6.3. The normalization, however, does not affect our estimates of the inclusive values V_{jtd}^{park}. From Equation

4, the inclusive value is computed as the log-sum of exponentiated mean utilities $\delta_{k|jtd}$ divided by β_{park}. Since the parameters and mean utilities are already normalized by β_{park}, we can consistently estimate V_{jtd}^{park} even before recovering β_{park}.

6.3 Estimation of Destination Preferences

We estimate destination preferences from the destination visit data described in Section 3 and the estimated V_{jtd}^{park} obtained from Section 6.2.

Since we observe information on visits to San Francisco by Bay Area residents, including each individual's home CBG, trip purpose, timestamp, and destination, we can estimate Equation 6 as a standard multinomial logit model. We define a market as a tuple of home CBG (h), trip purpose (c), and time (timeband t on date d). In each market, drivers choose which destination block j (of category c) to visit.

Identification. We instrument for V_{jtd}^{park} to address its correlation with the unobservables ω_{jctd}. Destinations with attractive unobserved characteristics draw more visitors. From Equation 14, higher destination visits translate into higher parking demand at nearby blocks, making parking more crowded and triggering demand-based price adjustments. This creates endogeneity between V_{jtd}^{park} and ω_{jctd}.

For each destination j, we use the variation in *moderately distant special events* as instruments for V_{jtd}^{park}. Note that these events are distinct from the *nearby special events* defined in Equation 6: while nearby events (e.g., within 400 meters of j) directly affect foot traffic to j, moderately distant events (e.g., 400 to 1000 meters from j) do not. Instead, these moderately distant special events attract crowds and increase parking demand at their nearby parking blocks, reducing the blocks' availability. This, however, can spill over and affect the parking options available to drivers visiting destination j.[30] This shifts V_{jtd}^{park} without directly entering drivers' utility for visiting j. In our main analysis, we define *moderately distant* as events occurring between 400 and 1000 meters from destination j, where the lower bound ensures the event does not directly affect j's attractiveness, and the upper bound ensures meaningful parking spillovers. Our results are robust to a wide range of distances.

Parameter Estimates. We estimate destination preferences using IV-2SLS. Table A.1 reports the first-stage results. Column (1) of Table 4 reports the estimates of $\boldsymbol{\beta}$. Drivers prefer destinations closer to home and favor street blocks with nearby special events. They also like POIs of all types and especially branded POIs that align with their trip purposes. Drivers exhibit considerable

[30]For example, an event occurring 500 meters from j can increase parking demand at some blocks 300 meters from j, which are within the parking choice set of j.

parking considerations. The nesting parameter β_{park} falls within the theoretically required range of zero to one, a restriction we do not impose in our estimation, which indicates that the assumed nested logit structure is not rejected by our estimates.

We also report the mean elasticities corresponding to the destination and parking estimates in Columns (1) and (2) of Table 5, respectively. In Stage 1, distance from home to destination has the highest elasticity among all attributes (mean elasticity $= 3.233$), implying that drivers are most responsive to destination proximity when making travel decisions. The mean elasticity of parking value, while lower than that of distance, is substantial relative to other physical attributes of destination j (such as nearby special events and POI composition), confirming the intuition that parking is a subsequent yet significant component of the travel decision. In Stage 2, demand for parking is elastic with respect to both expected availability and prices. This suggests that pricing can meaningfully influence parking choices, as drivers respond to parking fees both directly and indirectly through their effect on expected availability. We further examine the impacts of parking pricing in Section 8. The mean elasticity of steep terrain indicator is 0 because the variable's coefficient estimate is statistically insignificant.

7 The Welfare of Parking

In this section, we evaluate the welfare of parking using the framework in Section 5 and the estimates in Section 6. Understanding how parking availability and pricing affect driver surplus and cruising externalities is essential for evaluating alternative parking policies, which we discuss in Section 8.

7.1 Driver surplus

We begin by computing the driver surplus generated by each parking block. First, we calculate each driver's total trip value using the nested logit consumer surplus formula, converting surplus to monetary units by dividing by the price coefficient[31]

$$CS_{ictd} = \frac{1}{-\alpha_{price}} \log \left[\sum_{j} \exp(\delta_{jh(i)ctd}) \right] \tag{16}$$

where $\delta_{j|h(i)ctd}$ is defined in Equation 6.

Since drivers are heterogeneous with respect to their home CBGs and trip purposes, we com-

[31]Since prices interact with expected availability in our parking model, we interpret this as the expected dollar-equivalent value.

Table 4: Coefficient Estimates for Driver Choice Model

Covariate	Estimates: Destination ($\boldsymbol{\beta}$) (1)	Estimates: Parking ($\tilde{\boldsymbol{\alpha}}$) (2)
No. public-use garage		0.020
		(0.002)
Distance to parking		-3.338
		(0.192)
Expected availability $\times$ Hourly price		-0.902
		(0.041)
Expected availability $\times$ Street orientation		-0.065
		(0.020)
Expected availability $\times$ Steep terrain indicator		-0.009
		(0.029)
Distance to destination	-0.090	
	(0.003)	
Value of nearby parking	0.704	
	(0.064)	
No. special event	0.419	
	(0.042)	
No. branded establishment	0.200	
	(0.004)	
No. restaurant	0.016	
	(0.001)	
No. shop	0.023	
	(0.002)	
No. entertainment venue	0.009	
	(0.003)	
Home CBG FE	Yes	No
Category FE	Yes	No
Timeband FE	Yes	Yes
Neighborhood FE	No	Yes
N	4,078,850	41,210

Notes: This table reports estimates for the two-stage driver choice model. In particular, we report estimates for the two sets of parameters $\tilde{\boldsymbol{\alpha}}$ (parking choice) and $\boldsymbol{\beta}$ (destination choice) as discussed in Sections 6.2 and 6.3, respectively. Columns (1) and (2) show coefficient estimates for destination and parking choices, respectively. Standard errors are robust and reported in parentheses. Distance to parking (from destination) is normalized by 400 meters, the radius of the parking choice set around each destination. Branded establishments are defined as establishments in the same category as drivers' trip purposes that are affiliated with a chain. (e.g., branded restaurants for dining trips).

Table 5: Mean Elasticities for Driver Choice Model

Variable	Mean Elasticity: Destination (1)	Mean Elasticity: Parking (2)
No. public-use garage		0.013
Distance to parking		3.375
Expected availability		2.441
Hourly price		3.599
Street orientation		0.046
Steep terrain indicator		0
Distance to destination	3.233	
Value of nearby parking	0.529	
No. special event	0.009	
No. branded establishment	0.129	
No. restaurant	0.061	
No. shop	0.062	
No. entertainment venue	0.008	
N	4,078,850	41,210

Notes: This table reports mean elasticities corresponding to the coefficient estimates in Table 4. Columns (1) and (2) report the corresponding mean elasticities of destination and parking parameters, respectively. Distance to parking (from destination) is normalized by 400 meters, the radius of the parking choice set around each destination. Branded establishments are defined as establishments in the same category as drivers' trip purposes that are affiliated with a chain. (e.g., branded restaurants for dining trips). The mean elasticity of steep terrain indicator is 0 because the variable's coefficient estimate is statistically insignificant.

pute total travel value as a weighted sum across drivers

$$\begin{aligned} CS_{td} &= N_{td} \sum_{i,c,h} w_{ictd} CS_{ictd} \\ &= N_{td} \sum_{i,c,h} \underbrace{\frac{1}{N_{h(i)td}} s_{c|htd} s_{htd}}_{w_{ictd}} CS_{ictd} \end{aligned} \tag{17}$$

where N_{td} is the number of Bay Area drivers going to San Francisco during timeband t on date d, $N_{h(i)td}$ is the number of drivers from home CBG $h(i)$, and $s_{c|htd}$ and s_{htd} are defined in Appendix B.1.

In our framework, driver surplus from parking at block k during timeband t on date d is measured as the compensating variation from removing block k from the choice set of all parking locations in the commercial areas at that time. We have

$$CS_{ktd} = CS_{td} - CS_{td}^{-k} \tag{18}$$

where CS_{td}^{-k} represents the total travel value of all drivers when metered block k is not accessible for parking.

We then sum city revenue and driver surplus to compute each metered block's parking surplus and compare its PDV with the block's underlying land value to assess the economic return on parking use. This extends our discussion in Section 3, which compares only revenue with underlying land value. Figure 5 plots the PDV of parking surplus against land value for all metered blocks in the inner city. In both Figures 4 and 5, most blocks fall well below the 45-degree line, indicating that parking surplus captures only a small share of assessed land value. However, while in Figure 4, all blocks' revenues fall below the line, in Figure 5, 5.5% of the blocks lie above it, suggesting that metered parking can generate substantial surplus for certain high-value locations, exceeding the value of land uses there. On average, each block's PDV of parking surplus is over $700,000, accounting for roughly 40% of its underlying land value. By comparison, the PDV of revenue alone accounts for only about 23% of land value, as discussed in Section 4.

7.2 Cruising externalities

In our framework, we define a cruising trip as one where a driver attempts to park at a metered block in the inner city but finds it unavailable and resorts to a public-use garage instead (see Section 5.1). Although this is a simplification, it captures the essence of cruising for parking: when drivers cannot find a metered space, they move between parking locations, creating extra car trips. Efficient parking management can mitigate this externality.

Figure 5: Parking Surplus vs. Land Value

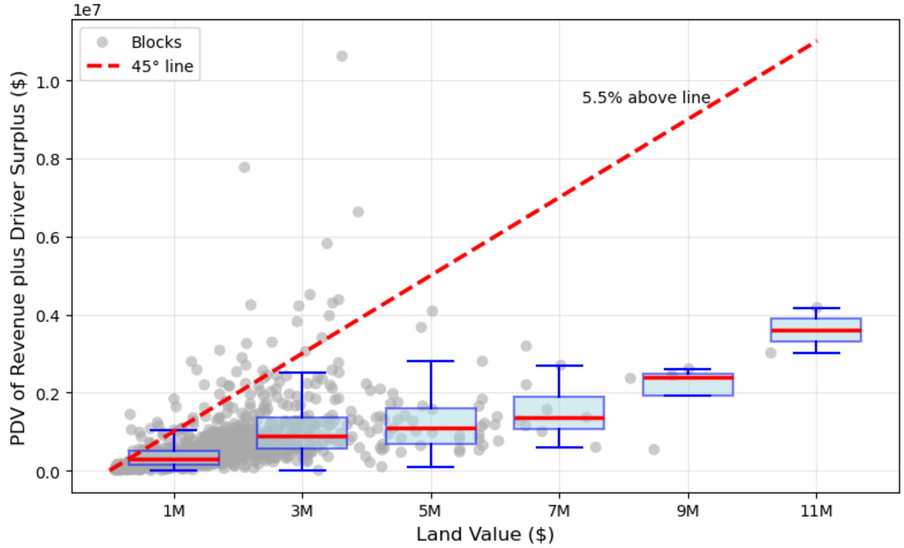

Notes: This figure presents the relationship between parking surplus, which is the sum of city revenue and driver surplus, and land value for all metered blocks in San Francisco's commercial areas. The horizontal axis shows the estimated land value for each parking block, calculated as the neighborhood land value per square foot multiplied by the standard metered stall size and the number of metered spaces on the block. The vertical axis shows the PDV of annual parking surplus calculated using a 7% discount rate. Each gray dot represents a parking block, and the red dashed line denotes the 45-degree line, where parking surplus equals land value. The box plots show the median and interquartile range of parking revenue within each land value bin. There are only 5.5% of metered blocks in the inner city whose PDV of parking surplus exceed their underlying land values.

The number of cruising trips at block k during timeband t on date d is therefore equal to the number of drivers attempting to park at k, which is estimated in Section 6.2, minus the number of successful parking transactions, which we observe in the SFMTA data. We then aggregate cruising trips across blocks to obtain a total of 68,834 cruising trips in the inner city during June 2019.

We compute the dollar costs of the cruising trips. Based on Weinberger et al. (2020), the average cruising time in San Francisco is 2 to 3 minutes per trip, while Cookson and Pishue (2017) estimates an average of 12 minutes spent searching for street parking per trip. Following the U.S. Department of Transportation's Guidance on Valuation of Travel Time in Economic Analysis, we measure the value of travel time savings (VTTS) as 50% of median household income per hour for local travel and 70% for intercity travel. Using the median household income in the Bay Area in 2019, these values are \$72.95 and \$102.14 per hour, respectively. We compute the cost of cruising as

$$\textit{Cost of Cruising} = \textit{Number of Cruising Trips} \times \textit{Average Cruising Time per Trip} \times \textit{VTTS}$$

With 68,834 cruising trips in the inner city in June 2019, this yields a monthly cruising cost ranging from approximately $200,000 (using 2.5 minutes cruising time and local VTTS) to $1,400,000 (using 12 minutes cruising time and intercity VTTS).

8 Counterfactual Analysis

In this section, we conduct two sets of counterfactuals to evaluate the trade-offs of alternative parking policies. First, we compare demand-based pricing with uniform pricing, the pricing scheme that most cities in the U.S. use. We set the uniform price at the level that maximizes total revenue. Second, we evaluate metered parking policies with varying parking supply and price adjustments. We consider the impacts of the counterfactual policies on parking revenue, driver surplus, and time costs of cruising. Under counterfactual prices and supply levels, we recompute equilibrium (fixed point) EA_{ktd} for every parking block across all dates and timebands in the data. See Appendix C for details on how to compute the new equilibrium EA_{ktd}.

8.1 Demand-based Pricing and Uniform Pricing

To compare demand-based and uniform pricing, we evaluate San Francisco's current demand-based pricing scheme against a uniform pricing scheme of $2.75/hour at every block, the uniform price level that maximizes revenue. We examine the two schemes in terms of parking revenue, driver surplus, and cruising time costs.

Table 6: Comparison of Demand-Based and Uniform Pricing

Inner city, per month	**Revenue** (1)	**Driver Surplus** (2)	**Cruising Trips** (3)	**Cruising Costs** (4)	**Total Welfare** (5)
Demand-based *(status quo)*	$2,092,918	$4,187,410	68,834	$251,086	$6,029,242
Uniform	$1,628,436	$4,688,966	224,123	$817,537	$5,499,865
Difference	$464,482	−$501,556	−155,289	−$566,451	$529,377

Notes: This table compares parking outcomes under demand-based pricing (status quo) and uniform pricing (a counterfactual where all blocks charge $2.75 per hour, the uniform price level that maximizes revenue). All values are monthly aggregates for San Francisco's commercial areas in June 2019. Column (1) is the total metered parking revenue in the commercial areas. Column (2) is total driver surplus from metered parking. Column (3) is the estimated number of trips where drivers attempt to park at a metered block but find it unavailable and resort to a public-use garage. Column (4) is the total time costs of cruising, estimated with the approach in Section 7.2. Column (5) is parking welfare, the sum of city revenue and driver surplus minus time costs of cruising. The "Difference" row shows demand-based minus uniform pricing outcomes.

Table 6 presents the comparison of demand-based and uniform pricing schemes. On the one hand, demand-based pricing generates higher revenue than the revenue-maximizing uniform price:

the city collects approximately $464,000 more in monthly parking revenue from commercial areas. On the other hand, drivers enjoy higher surplus under the uniform pricing scheme, which offsets the revenue gains. These results arise because, although demand-based pricing allocates parking more efficiently, drivers face substantially higher prices overall. As of 2019, roughly half of the blocks had prices above $2.75, and the highest rate was nearly $8, about three times the uniform rate of $2.75. This pattern is consistent with the elasticity estimates reported in Table 5, which shows that at the actual 2019 price levels, drivers on average are more sensitive to price changes than to changes in expected availability.

The criterion with the most significant improvement is cruising. Demand-based pricing reduces cruising trips by approximately 155,000 per month, a 69% reduction. This comes from the improved availability of crowded parking blocks under demand-based pricing relative to uniform pricing. Using the valuation approach from Section 7.2, this translates to a reduction in cruising costs of approximately $500,000 to $3,000,000 per month, depending on assumptions about cruising time and VTTS. Overall, under the status quo demand-based pricing, curbside parking welfare, measured as the sum of city revenue and driver surplus minus cruising costs, increases by nearly 10% relative to revenue-maximizing uniform pricing.

8.2 Trade-offs of Alternative Parking Policies

We evaluate two types of policy interventions. First, we consider pricing adjustments. Given that our sample has 886 parking blocks, two timebands (morning and afternoon), and two day types (weekday and weekend), which yields 3,544 single prices to sets, finding the optimal prices is computationally challenging. We resort to a simpler pricing rule under which prices are set according to

$$P_{ktd}^{new} = \max(0, P_{ktd}^{current} + 0.25 \times \lambda) \tag{19}$$

where P_{ktd}^{new} denotes the new price at block k during timeband t on day type d (weekday or weekend), $P_{ktd}^{current}$ is the current price at the corresponding block and time, and λ is an integer adjustment factor that can take negative, zero, or positive values. For example, $\lambda = 2$ corresponds to a $0.50 increase, while $\lambda = -2$ represents a $0.50 decrease. This rule applies a common adjustment to prices across all blocks, timebands, and day types. The pricing adjustment is consistent with San Francisco's most recent intervention, which implements a citywide $0.25 increase.

Second, we evaluate capacity reductions by simulating scenarios where the number of metered parking spaces per block is reduced to 25%, 50%, or 75% of baseline capacity. These counterfactuals capture the effects on parking of policies that reduce metered parking supply and potentially redirect curb space to alternative uses.[32] For each capacity level, we also examine how parking availability and parking outcomes respond to price adjustments, characterizing the trade-offs among alternative curbside parking policies.

Figure 6 plots parking availability against price changes at baseline capacity. As shown, higher prices improve availability, increasing the share of blocks with high availability (i.e., those available more than 75% of the time). Figure 7 examines how citywide price adjustments affect parking welfare, including city revenue, driver surplus, cruising costs, and total welfare, at the current capacity. In Panel (a), revenue follows an inverted U-shaped curve, with the current pricing scheme (marked with an X) positioned near the revenue-maximizing point within the pricing family defined by Equation 19. Moving to the optimum, however, would increase monthly revenue by only less than 1%. In Panel (b), driver surplus goes up as prices go down, yet the rate of increase slows down as prices drop further. This is because, at current prices, parking availability is generally high (as shown in Figure 6), making consumers relatively more responsive to price changes. However, as prices fall further, equilibrium effects become more pronounced, making availability decline faster and consumers less responsive to prices. In Panel (c), cruising costs go down as prices increase, which aligns with improved availability. Finally, taking all factors into account, Panel (d) shows that parking welfare exhibits a hump-shaped relationship with price changes. This suggests that at the status quo, reducing prices is the dominant strategy for improving parking outcomes. If the local government prioritizes revenue, prices should be lowered slightly by $0.25/hour. If the government instead emphasizes parking welfare, prices should be reduced by $2/hour to achieve the maximum welfare.

Figure 8 examines how parking welfare and revenue respond to different combinations of parking supply reductions and price adjustments. Panel (a) shows total parking welfare across four capacity scenarios (25%, 50%, 75%, and 100% of current capacity) and various price changes, while Panel (b) shows the revenue implications. Several patterns emerge. First, reducing capacity alone, without price changes, results in only a modest change in parking revenue but a substantial decline in parking welfare. Second, at the status quo capacity, both the value-maximizing and revenue-maximizing prices require reducing prices. In contrast, when capacity becomes more constrained, the city needs to raise prices to achieve welfare and revenue maximization. Under extreme capacity reductions, availability becomes so limited that drivers no longer benefit from price cuts.

[32]There are caveats to our analysis. First, we do not model how conversion of curb space to alternative uses affects foot traffic. However, we expect this effect to be marginal in the short to medium run, as curb space conversion and the resulting changes in area attractiveness take time to materialize. Second, due to data limitations, our framework does not model the changes in mode choice in response to parking policies. This is less concerning for short- to medium-run policy analysis, as mode switching and vehicle ownership decisions tend to adjust more slowly than parking and destination choices. Given the limited research on the relationship between mode choice and parking policies (Krishnamurthy and Ngo, 2020), understanding how mode choice responds to parking policies remains an interesting and important direction for future research on parking interventions.

Figure 6: Shares of Block with High Availability (> 75%) under Counterfactual Pricing

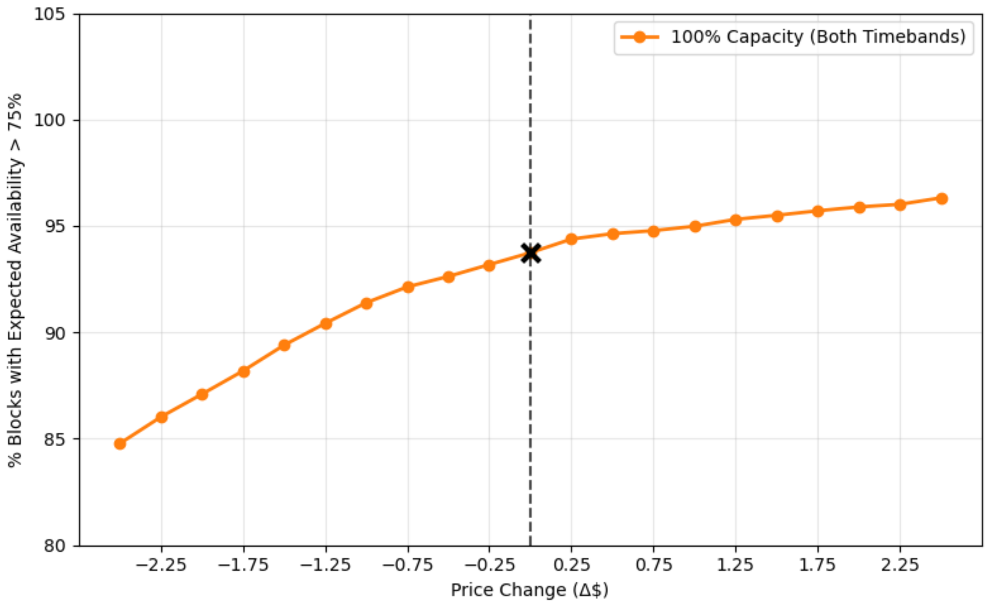

Notes: This figure shows how citywide price adjustments affect the share of blocks with high availability in San Francisco's commercial areas, defined as blocks where the probability of finding an open spot during a timeband exceeds 75%. The horizontal axis shows the total price change in dollars, where price changes (Δ$) follow the pricing rule in Equation 19. The X marks the current pricing scheme (no price change). All outcomes are averaged over timebands and aggregated over dates in June 2019.

Figure 9 plots parking welfare across different capacity levels under a $1.25 citywide price reduction, interpolating between simulated scenarios. Relative to the status quo, reducing parking capacity by about 6%, combined with a $1.25 price reduction, yields a similar level of parking welfare as the current system, with only a modest decline in city revenue. Lowering prices compensates drivers for the decline in availability due to reduced supply. Among all price adjustments we consider, a $1.25 reduction allows the largest feasible capacity reduction while holding parking welfare at its current level. Figure A.2 shows all combinations of capacity and price adjustments that can achieve the status quo welfare. Overall, coordinating supply reductions with strategic price adjustments can help local governments balance alternative policy objectives.

9 Conclusion

This paper quantifies the welfare effects of curbside parking, including city revenue, driver surplus, and time costs of cruising, and uses these measures to evaluate alternative policies for managing curb space through parking instruments. We develop a structural model of drivers' joint destination and parking choices to estimate parking preferences and the extent to which parking

Figure 7: Revenue, Driver Surplus, Cruising Costs, and Welfare under Counterfactual Pricing

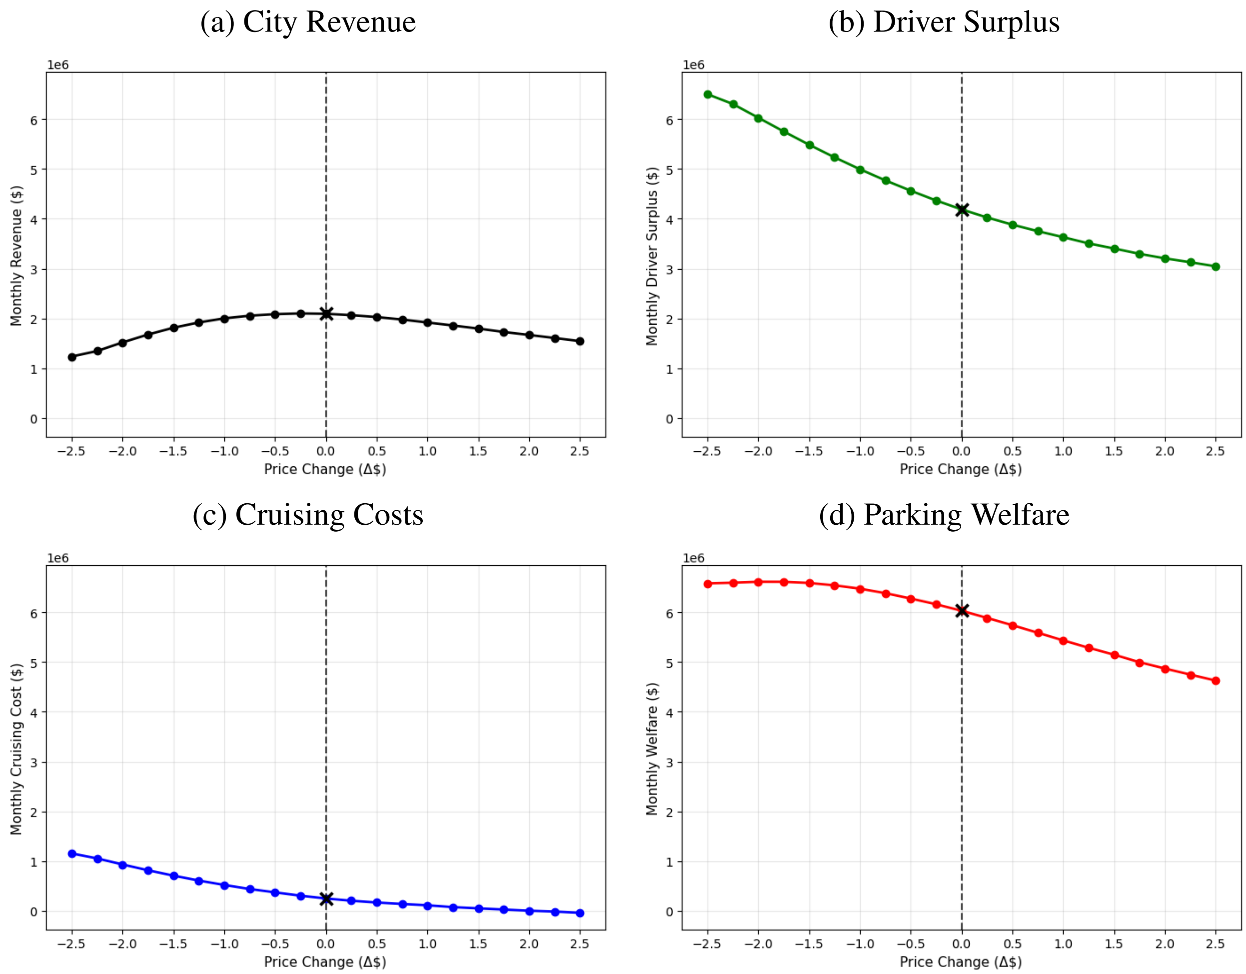

Notes: This figure shows how citywide price adjustments affect parking outcomes in San Francisco's commercial areas. Each panel plots parking outcomes against price changes ($\Delta\$$) following the pricing rule in Equation 19. The horizontal axis shows the total price change in dollars. Panel (a) shows monthly parking revenue, Panel (b) shows driver surplus, Panel (c) shows cruising costs, and Panel (d) shows parking welfare, measured as the sum of city revenue and driver surplus minus cruising costs. The X marks the current pricing scheme (no price change). All outcomes are averaged over timebands and aggregated over dates in June 2019.

considerations factor into travel decisions. Our estimates show that in general, parking is subsequent yet has a meaningful influence on drivers' travel decisions. Drivers respond to both parking prices and expected availability, making pricing an effective tool for managing curbside parking demand.

Using the model, we quantify driver surplus. We then compare each block's PDV of parking surplus, the sum of city revenue and driver surplus, with its underlying land value. We find that while the PDV of revenue alone is about 20% of land value, the PDV of city revenue plus driver surplus amounts to roughly 40% of land value. Moreover, 5.5% of blocks have a PDV of parking surplus that exceeds land value, compared to no blocks when considering revenue alone.

Figure 8: Parking Welfare and Revenue under Counterfactual Pricing and Supply

(a) Parking Welfare

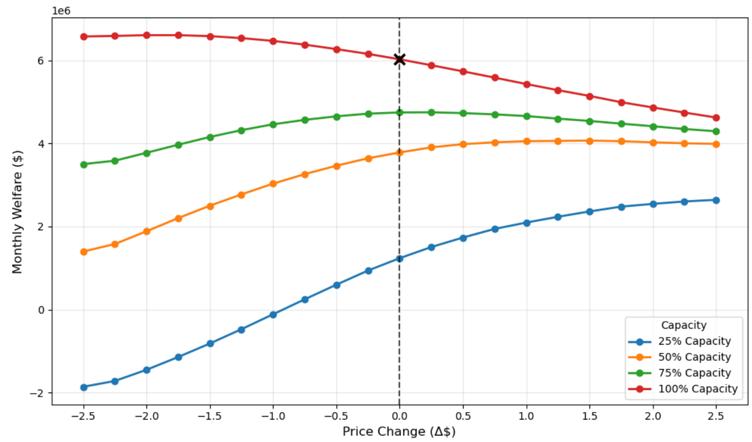

(b) City Revenue

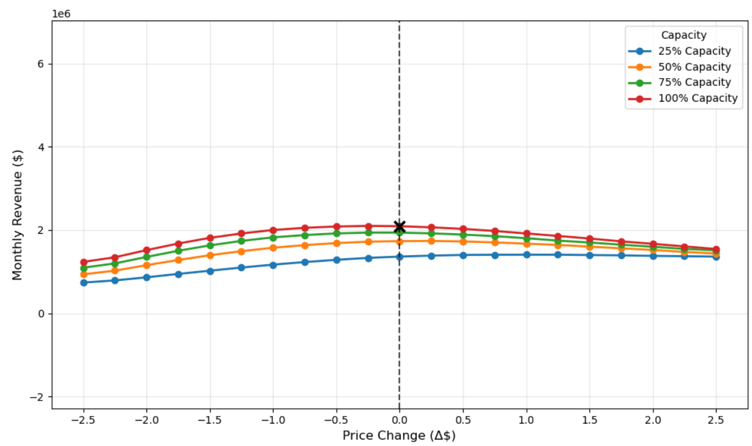

Notes: This figure shows how parking outcomes respond to combinations of capacity reductions and citywide price adjustments. Both panels plot parking outcomes against price changes ($\Delta\$$) following the pricing rule in Equation 19. The horizontal axis shows the total price change in dollars. The four colored lines represent different capacity scenarios: 100% (current capacity), 75%, 50%, and 25% of current capacity, applied uniformly across all parking blocks and times. Panel (a) shows monthly total parking welfare (revenue plus driver surplus minus cruising costs). Panel (b) shows monthly parking revenue. The X marks the current pricing scheme (no price change). All outcomes are averaged over timebands and aggregated over dates in June 2019.

Our counterfactual exercise shows that, compared to a revenue-maximizing uniform pricing scheme, San Francisco's demand-based pricing generates about 30% more revenue while reducing cruising trips by nearly 70%. Another set of counterfactual analysis suggests that reducing parking supply by roughly 6% while reducing the status quo demand-based prices by \$1.25 citywide preserves parking welfare, with only a modest revenue loss. This implies that cities can achieve more flexible use of curb space through coordinated parking pricing and curb space provision policies.

Figure 9: Welfare under Counterfactual Capacity and $1.25 Price Reduction

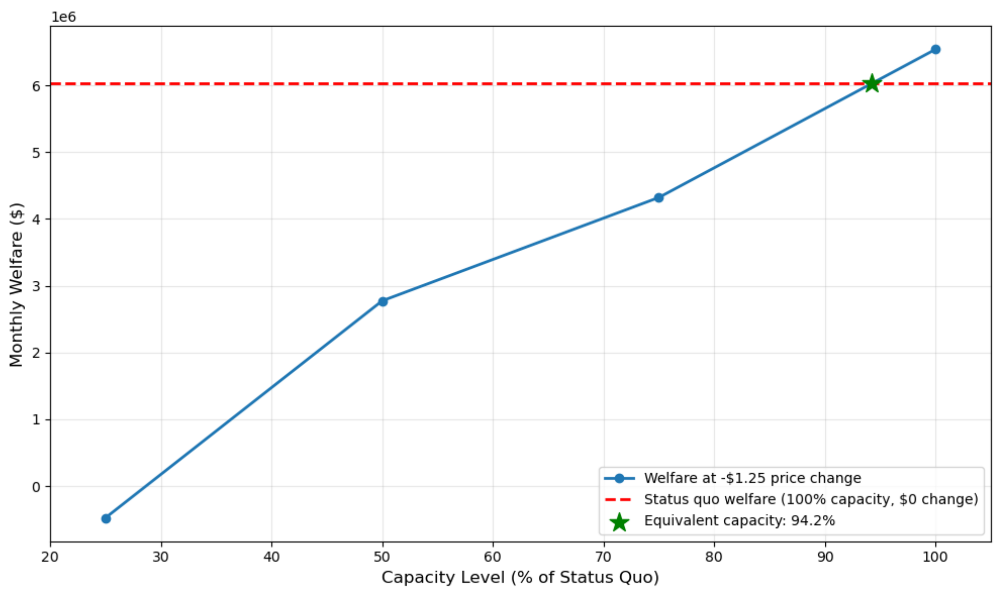

Notes: This figure shows monthly parking welfare under different capacity levels combined with a $1.25 price reduction across all blocks. The horizontal axis shows parking supply (capacity) as a percentage of the current level. The vertical axis shows parking welfare (city revenue plus driver surplus minus cruising costs). The blue line shows welfare at each capacity level with prices reduced by $1.25. The red dashed line shows baseline welfare under current pricing and capacity. The green star marks the capacity level (94.2%) where $1.25 price reductions achieve the same welfare as the status quo, identified through interpolation. All outcomes are averaged over timebands and aggregated over dates in June 2019.

References

Almagro, M., Barbieri, F., Castillo, J. C., Hickok, N. G., & Salz, T. (2024). *Optimal urban transportation policy: Evidence from Chicago* (NBER Working Papers No. 32185). National Bureau of Economic Research.

Almagro, M., & Domínguez-Iino, T. (2025). Location sorting and endogenous amenities: Evidence from Amsterdam. *Econometrica*, *93*(3), 1031–1071.

Anderson, S. P., & De Palma, A. (2004). The economics of pricing parking. *Journal of Urban Economics*, *55*(1), 1–20.

Arnott, R., & Inci, E. (2006). An integrated model of downtown parking and traffic congestion. *Journal of Urban Economics*, *60*(3), 418–442.

Arnott, R., & Inci, E. (2010). The stability of downtown parking and traffic congestion. *Journal of Urban Economics*, *68*(3), 260–276.

Arnott, R., Inci, E., & Rowse, J. (2015). Downtown curbside parking capacity. *Journal of Urban Economics*, *86*, 83–97.

Arnott, R., & Rowse, J. (1999). Modeling parking. *Journal of Urban Economics*, *45*(1), 97–124.

Athey, S., Blei, D., Donnelly, R., Ruiz, F., & Schmidt, T. (2018). Estimating heterogeneous consumer preferences for restaurants and travel time using mobile location data. *AEA Papers and Proceedings*, *108*, 64–67.

Azar, J. A., Berry, S. T., & Marinescu, I. (2022). *Estimating labor market power* (NBER Working Papers No. 30365). National Bureau of Economic Research.

Barwick, P. J., Li, S., Waxman, A., Wu, J., & Xia, T. (2024). Efficiency and equity impacts of urban transportation policies with equilibrium sorting. *American Economic Review*, *114*(10), 3161–3205.

Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, *63*(4), 841–890.

Buchholz, N. (2022). Spatial equilibrium, search frictions, and dynamic efficiency in the taxi industry. *The Review of Economic Studies*, *89*(2), 556–591.

Cao, Y., Chevalier, J. A., Handbury, J., Parsley, H., & Williams, K. R. (2024). *Distributional impacts of the changing retail landscape* (Cowles Foundation Discussion Papers No. 2415). Cowles Foundation for Research in Economics, Yale University.

Cardell, N. S. (1997). Variance components structures for the extreme-value and logistic distributions with application to models of heterogeneity. *Econometric Theory*, *13*(2), 185–213.

Castillo, J. C. (2025). Who benefits from surge pricing? *Econometrica*, *93*(5), 1811–1854.

Chaniotakis, E., & Pel, A. J. (2015). Drivers' parking location choice under uncertain parking availability and search times: A stated preference experiment. *Transportation Research Part A: Policy and Practice*, *82*, 228–239.

Chatman, D. G., & Manville, M. (2014). Theory versus implementation in congestion-priced parking: An evaluation of SFpark, 2011–2012. *Research in Transportation Economics*, *44*, 52–60.

Conlon, C., & Gortmaker, J. (2020). Best practices for differentiated products demand estimation with PyBLP. *The RAND Journal of Economics*, *51*(4), 1108–1161.

Cook, C. (2024). Heterogeneous preferences for neighborhood amenities: Evidence from GPS data. *The Review of Economics and Statistics*, 1–43.

Cook, C., Kreidieh, A., Vasserman, S., Allcott, H., Arora, N., van Sambeek, F., Tomkins, A., & Turkel, E. (2025). *The short-run effects of congestion pricing in New York City* (NBER Working Papers No. 33584). National Bureau of Economic Research.

Cook, C., & Li, P. Z. (2025). *Value pricing or lexus lanes? The distributional effects of dynamic tolling* [Working Paper].

Cookson, G., & Pishue, B. (2017). *The impact of parking pain in the US, UK and Germany* (Research Report). INRIX Research.

Couture, V. (2016). *Valuing the consumption benefits of urban density* [Working Paper].

Couture, V., Dingel, J. I., Green, A. E., & Handbury, J. (2025). *Demographic preferences and income segregation* (NBER Working Papers No. 33386). National Bureau of Economic Research.

Dalla Chiara, G., & Goodchild, A. (2020). Do commercial vehicles cruise for parking? Empirical evidence from Seattle. *Transport Policy*, *97*, 26–36.

Davis, D. R., Dingel, J. I., Monras, J., & Morales, E. (2019). How segregated is urban consumption? *Journal of Political Economy*, *127*(4), 1684–1738.

Durrmeyer, I., & Martínez, N. (2024). *The welfare consequences of urban traffic regulations* [Working Paper].

Erlang, A. K. (1909). The theory of probabilities and telephone conversations. *Nyt Tidsskrift for Matematik. Serie B*, *20*, 33–39.

Feldman, P., Li, J., & Tsai, H. T. (2022). Welfare implications of congestion pricing: Evidence from SFpark. *Manufacturing & Service Operations Management*, *24*(2), 1091–1109.

Frechette, G. R., Lizzeri, A., & Salz, T. (2019). Frictions in a competitive, regulated market: Evidence from taxis. *American Economic Review*, *109*(8), 2954–2992.

Grabar, H. (2024). *Paved paradise: How parking explains the world*. Penguin.

Gragera, A., Hybel, J., Madsen, E., & Mulalic, I. (2021). A model for estimation of the demand for on-street parking. *Economics of Transportation, 28*, 100231.

Gupta, A., Van Nieuwerburgh, S., & Kontokosta, C. (2022). Take the Q train: Value capture of public infrastructure projects. *Journal of Urban Economics*, *129*, 103422.

Kashner, Z., & Ross, B. (2025). *Measuring and mitigating traffic externalities* [Working Paper].

Kaufman, M., Formanack, M., Gray, J., & Weinberger, R. (2012). *Contemporary approaches to parking pricing: A primer* (FHWA Report No. FHWAHOP12026). Federal Highway Administration.

Kelly, J. A., & Clinch, J. P. (2009). Temporal variance of revealed preference on-street parking price elasticity. *Transport Policy*, *16*(4), 193–199.

Krishnamurthy, C. K. B., & Ngo, N. S. (2020). The effects of smart-parking on transit and traffic: Evidence from SFpark. *Journal of Environmental Economics and Management*, *99*, 102273.

Lehner, S., & Peer, S. (2019). The price elasticity of parking: A meta-analysis. *Transportation Research Part A: Policy and Practice*, *121*, 177–191.

McFadden, D. (1973). Conditional logit analysis of qualitative choice behavior. In *Frontiers in econometrics* (pp. 105–142).

McFadden, D. (1978). Modelling the choice of residential location. In A. Karlqvist, L. Lundqvist, F. Snickars, & J. Weibull (Eds.), *Spatial interaction theory and planning models* (pp. 75–96). North-Holland.

Millard-Ball, A., Weinberger, R. R., & Hampshire, R. C. (2014). Is the curb 80% full or 20% empty? Assessing the impacts of San Francisco's parking pricing experiment. *Transportation Research Part A: Policy and Practice*, *63*, 76–92.

Millard-Ball, A., Hampshire, R. C., & Weinberger, R. (2020). Parking behaviour: The curious lack of cruising for parking in San Francisco. *Land Use Policy*, *91*, 103918.

Miyauchi, Y., Nakajima, K., & Redding, S. J. (2025). The economics of spatial mobility: Theory and evidence using smartphone data. *The Quarterly Journal of Economics*, *140*(4), 2507–2570.

Ostermeijer, F., Koster, H., Nunes, L., & van Ommeren, J. (2022). Citywide parking policy and traffic: Evidence from Amsterdam. *Journal of Urban Economics*, *128*, 103418.

Pierce, G., & Shoup, D. (2013). Getting the prices right: An evaluation of pricing parking by demand in San Francisco. *Journal of the American Planning Association*, *79*(1), 67–81.

Pierce, G., Willson, R., & Shoup, D. (2015). Optimizing the use of public garages: Pricing parking by demand. *Transport Policy*, *44*, 89–95.

Redding, S. J., & Turner, M. A. (2015). Transportation costs and the spatial organization of economic activity. In *Handbook of regional and urban economics* (pp. 1339–1398, Vol. 5). Elsevier.

Rosaia, N. (2020). *Competing platforms and transport equilibrium: Evidence from New York City* [Working Paper].

SFMTA. (2014a, June). *SFpark pilot project evaluation* (Technical Report). San Francisco Municipal Transportation Agency.

SFMTA. (2014b, June). *SFpark pilot project evaluation summary* (Technical Report). San Francisco Municipal Transportation Agency.

Shoup, D. (2005). *High cost of free parking*. Routledge.

Shoup, D. (2018). *Parking and the city*. Routledge.

Su, Y. (2022). Measuring the value of urban consumption amenities: A time-use approach. *Journal of Urban Economics*, *132*, 103495.

Train, K. E. (2009). *Discrete choice methods with simulation* (2nd ed.). Cambridge University Press.

Tsivanidis, N. (2023). *Evaluating the impact of urban transit infrastructure: Evidence from Bogotá's TransMilenio* [Working Paper].

Veraset. (2022). *Visits*. Dewey Data. https://doi.org/10.82551/TP7E-VA02

Vickrey, W. (1954). The economizing of curb parking space. *Traffic Engineering*, *25*(2), 62–67.

Weinberger, R. R., Millard-Ball, A., & Hampshire, R. C. (2020). Parking search caused congestion: Where's all the fuss? *Transportation Research Part C: Emerging Technologies*, *120*, 102781.

Appendices

A Supplemental Figures and Tables

Figure A.1: Parking Block versus Street Block

Notes: This map shows the distinction between parking blocks and street blocks in San Francisco. Yellow strips represent parking blocks, defined as the street segment between a pair of opposing blockfaces. Gray areas represent street blocks.

Figure A.2: Price–Capacity Adjustment Combination to Match Baseline Welfare

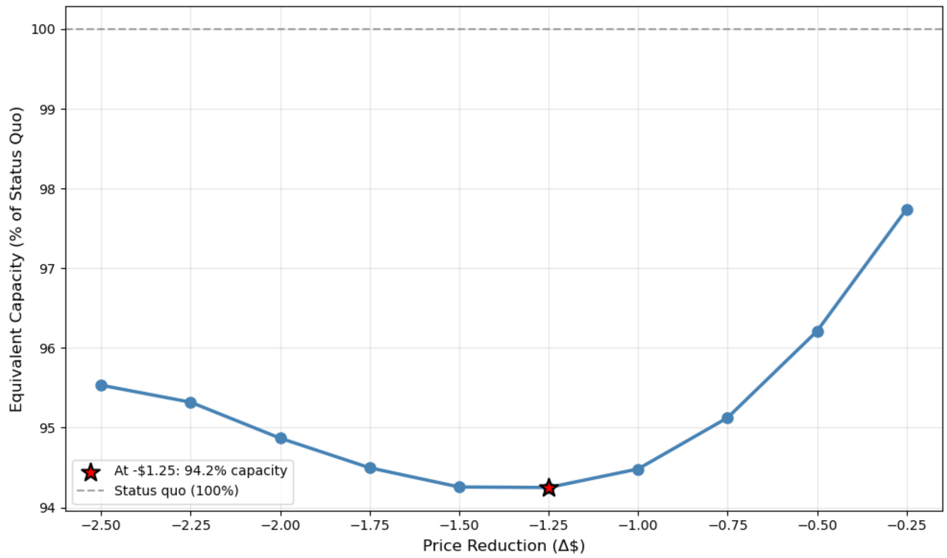

Notes: This figure shows alternative combinations of price and capacity adjustments that maintain baseline welfare (status quo at 100% capacity with no price change). The vertical axis shows different capacity levels (as a percentage of status quo). The horizontal axis shows corresponding price adjustments ranging from -$2.50 to -$0.25 to achieve baseline welfare. The dashed horizontal line indicates the status quo capacity level (100%). The red star marks the point of maximum capacity reduction: a $1.25 price reduction paired with 5.8% capacity reduction (to 94.2%) maintains baseline welfare. The U-shaped pattern indicates that moderate price reductions (around -$1.25) are most effective at compensating for capacity reductions, while larger price reductions require capacity levels closer to the status quo to maintain the same welfare level.

Table A.1: First Stage of Destination Choice Estimation

Covariate	Estimate	Std. Error
Distance to destination	0.033	0.000
No. special event	−0.412	0.003
No. branded establishment	−0.037	0.000
No. restaurant	0.008	0.000
No. shop	−0.012	0.000
No. entertainment venue	0.003	0.000
No. moderately distant event	−0.178	0.001
Home CBG FE	Yes	
Category FE	Yes	
Timeband FE	Yes	

Notes: This table reports coefficient estimates from the first stage of destination choice estimation. The dependent variable is the inclusive value of parking at each destination. We use *moderately distant* special events occurring between 400 and 1,000 meters of each destination to instrument for the endogenous inclusive parking value of that destination.

Table A.2: Binomial GLM Estimation with Logit Link for EA_{ktd}

Variable	Estimate	Std. Error
Baseline availability prediction	7.703	0.005
Non-metered parking closures	−0.103	0.003
Date FE	Yes	

Notes: This table reports coefficient estimates from the binomial GLM with a logit link for EA_{ktd}, as specified in Equations 10 and 12. We use Maximum Likelihood Estimation.

B Technical Appendix

B.1 Market Shares

The share of drivers visiting j conditional on coming from home CBG h, regardless of trip purpose, is

$$s_{j|htd} = \sum_{c} s_{j|hctd} \times s_{c|htd} \tag{B.1}$$

where $c \in \{\text{dining, shopping, entertainment}\}$, $s_{j|hctd}$ comes from Equation 7, and $s_{c|htd}$, the probability of drivers visiting category c from home CBG h during timeband t on date d, comes from our GPS data.

The share of drivers visiting j during timeband t on date d is then

$$s_{jtd} = \sum_{h} s_{j|htd} \times s_{htd} \tag{B.2}$$

where s_{htd}, the empirical distribution of drivers' home CBG, comes from our GPS data.

The distribution of foot traffic to destination $j \in \mathcal{J}/\{0\}$, conditional on visiting the inner city, is represented by

$$s_{jtd}^{inner} = \frac{s_{jtd}}{\sum_{j' \in \mathcal{J}/\{0\}} s_{j'td}} \tag{B.3}$$

where s_{jtd} is computed in Equation B.2.

The share of drivers attempting parking at k during timeband t on date d, conditional on visiting the inner city, is

$$s_{ktd} = \sum_{j \in \mathcal{J}/\{0\}} s_{k|jtd} \times s_{jtd}^{inner} \tag{B.4}$$

where $s_{k|jtd}$ comes from Equation 8.

B.2 Erlang B Formula

In queueing theory, the Erlang B formula represents the loss probability in systems where customers arrive and leave if all servers are busy. This probability is a function of the system's capacity (supply) and the offered load (demand), the product of the arrival rate and the mean service duration. The predicted availability A_{kt} is then computed as one minus the Erlang B probability

$$A_{kt}(\lambda_{kt}, Dur_{kt}, Cap_{kt}) = 1 - \left[\underbrace{\frac{(\lambda_{kt} Dur_{kt})^{Cap_{kt}}}{Cap_{kt}!} \Big/ \sum_{\iota=0}^{Cap_{kt}} \frac{(\lambda_{kt} Dur_{kt})^{\iota}}{\iota!}}_{\text{Erlang B}} \right] \tag{B.5}$$

where λ_{kt} is the average arrival rate to k during timeband t, Dur_{kt} is the mean service duration, and Cap_{kt} is the capacity of system k during timeband t.

In our parking context, we infer the mean duration of parking (Dur_{kt}) and the number of parking spaces per block (Cap_{kt}) from the parking transaction and inventory data, respectively.[33] The arrival rate λ_{kt} is computed as the average number of people attempting parking at k per timeband divided by the length of the timeband (180 minutes). See Section 6.2 for how to derive the number of people attempting parking at k from the metered parking transactions data.

[33] In this paper, we treat average duration as a fixed characteristic of a parking block. The duration of parking comes from the duration of an amenity trip, and thus, the average duration of parking at a block is highly correlated with the composition of nearby amenities, which is held fixed in our model. Previous work has found mixed findings on whether parking prices affect parking duration. A report by SFMTA (2014a) and Krishnamurthy and Ngo (2020) find evidence of shorter duration when prices increase, while Chatman and Manville (2014) finds no statistical association.

C Computation of New Equilibrium

Algorithm 1 Equilibrium Computation Under Counterfactual Prices and Supply

Given counterfactual prices and supply, initialize EA^0_{ktd} for all (k, t, d)
repeat
    Compute $s_{k|jtd}$ and V^{park}_{jtd} using $\boldsymbol{\alpha}$ (Sec. 6.2)
    Compute $s_{j|hctd}$ using V^{park}_{jtd} and $\boldsymbol{\beta}$ (Sec. 6.3)
    Compute s_{ktd} (Sec. B.1)
    Recover N_{ktd} and average across dates to obtain N_{kt}
    Compute arrival rate $\lambda_{kt} = N_{kt}/180$ (Sec. B.2)
    Compute A_{kt} given λ_{kt}, counterfactual supply, and observed mean parking duration
    Update expected availability EA^1_{ktd} using $\boldsymbol{\eta}$ (Sec. 6.1)
    Compute $\|EA^1_{ktd} - EA^0_{ktd}\|$
    Set $EA^0_{ktd} \leftarrow EA^1_{ktd}$
until convergence

Since the dimension of EA_{ktd} is large and our framework is highly non-linear, it is challenging to rigorously prove that the algorithm converges to a unique equilibrium. Instead, we examine this computationally. For each counterfactual scenario, we use multiple initial guesses for EA_{ktd}: we initialize all blocks' expected availability at uniformly low, medium, and high levels, and also use the observed EA_{ktd} as a starting point. The resulting equilibria are highly consistent across starting values, with an average difference of only 0.6% across all (k, t, d) observations.