Transportation Research Part A 193 (2025) 104389

Contents lists available at ScienceDirect

Transportation Research Part A

journal homepage: www.elsevier.com/locate/tra

Pricing parking for fairness — A simulation study based on an empirically calibrated model of parking behavior

Jakob Kappenberger [a,*], Heiner Stuckenschmidt [a], Frederic Gerdon [b]

[a] *Mannheim Center for Data Science, University of Mannheim, B6, 26, Mannheim, 68159, Germany*
[b] *Mannheim Centre for European Social Research, University of Mannheim, A5, 6, Mannheim, 68159, Germany*

ARTICLE INFO

Keywords:
Parking choice
Dynamic pricing
Agent-based simulation
Artificial intelligence
Machine Learning
Fairness

ABSTRACT

It has been widely recognized that public parking, if not managed correctly, can significantly decrease a city's quality of life due to increased traffic and its impact on mobility and the environment. To avoid these negative effects, various parking policies have been proposed to reduce traffic while guaranteeing high accessibility, especially in city centers. This work investigates different pricing policies for public parking, including dynamic pricing and Machine Learning-based strategies that can directly optimize policy goals, such as improving mobility or accessibility. In doing so, we pay special attention to an aspect often ignored when implementing pricing policies for public parking: fairness with regard to equal outcomes for different social groups. Since the effects of pricing policies are very sensitive to financial inequality, we specifically investigate the impact of policies on different income groups. As a foundation for these experiments, we introduce a parking simulation featuring an empirically calibrated behavioral model of parking. We find that (1) dynamic pricing schemes may negatively impact fairness; (2) fair pricing for parking may require different fees for individual social groups; (3) focusing on single policy goals when devising pricing for parking results in unintended consequences; (4) Machine Learning shows potential for creating pricing strategies combining different policy goals.

1. Introduction

The number of cars in many urban areas continues to rise. For instance, the median growth rate of the motorization rate in OECD countries between 2000 and 2020 amounted to 57%, with only one country exhibiting negative growth (OECD, 2023). Consequently, parking spaces grow increasingly scarce, resulting in more congestion due to a higher number of vehicles searching for parking simultaneously. This is frequently attributed to inefficient management of public parking spaces, as they are often priced far too cheaply in relation to the social costs incurred (Shoup, 2011). At the same time, due to the challenges presented by climate change, there is pressure on policymakers to reduce traffic in city centers and incentivize the use of other modes of transportation, such as bikes or public transport (Litman, 2019). Using pricing policies that adjust parking costs depending on different parameters in order to impact traffic flow is a popular tool for policymakers to address these challenges. Consequently, it has already been implemented in multiple agglomerations, such as Madrid or San Francisco (Friesen and Mingardo, 2020).

Responding to this development, scientific interest in the impact of these systems has grown in the last decade. Most of these studies have focused on how well these pricing schemes achieve the desired degree of utilization of the affected parking spaces or improve traditional performance metrics, such as traffic flow (e.g., Pierce and Shoup (2013)). However, as it is generally accepted

* Corresponding author.
E-mail address: jakob.kappenberger@uni-mannheim.de (J. Kappenberger).

https://doi.org/10.1016/j.tra.2025.104389

Available online 6 February 2025

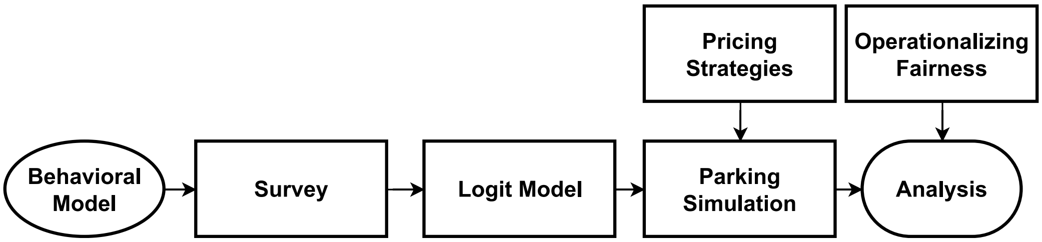

Fig. 1. Research workflow of this study, from defining a theoretical model of parking choice to analyzing the outcomes produced by the parking simulation using different pricing strategies with regard to fairness and other metrics.

that public parking is often too cheap, it is easily conceivable that the potentially more aggressive pricing conforming to aims, such as reducing overall traffic, disproportionately affects less affluent members of society. Nonetheless, the fairness aspect of dynamic pricing for parking, as well as parking in general, appears strikingly under-researched.

Following Shoup (2011, p. 299), the price of public parking should not be simply understood as a price set by the free market, but rather "a public price for a public service, and it should be set to achieve the public goals of improving transportation, land use, and the environment". Supplementing the approaches deployed in academic and municipal practice based on simple condition rules to achieve these goals, we propose a pricing system based on Machine Learning (ML) capable of directly optimizing various individual aims potentially pursued by municipalities. To test this system and analyze existing strategies, we design a behavioral model of parking that we calibrate based on data from a representative survey and use to extend and improve the parking simulation introduced in Kappenberger et al. (2022). We analyze the performance of the different pricing strategies more broadly by considering their impact on fairness and can thereby systematically investigate potential unintended consequences of focusing on individual aspects, such as traffic flow. In doing so, we follow an interdisciplinary approach, combining methods from transportation research, the social sciences, and artificial intelligence research. Fig. 1 illustrates our research workflow.[1]

In summary, our contributions are the following:

- **Theory**: In Section 2, we analyze the contributions made by related work in the space of parking models. Based on these considerations, we define a behavioral model of parking choice in Section 3.1. Finally, in Section 5, we conceptualize the fairness of parking as the equity of the outcomes obtained by different social groups and propose a group-based equity measure to analyze the degree of equity a pricing strategy achieves.
- **Models**: Based on a representative survey conducted in Germany, we estimate a logit regression model of parking choice (presented in Section 3.3), which we use to calibrate the behavioral model in our Agent-Based Model (ABM) for parking. Section 4 serves to introduce said simulation of priced parking modeled after the city of Mannheim, Germany. Furthermore, we design a ML-based pricing system for parking that we use to optimize a comprehensive set of policy goals (by specifying several reward functions) to maximize occupancy, fairness, and traffic flow.
- **Analysis**: We systematically examine the impact of the ML-based pricing system as well as of two baselines in our simulation and find that the ML-governed pricing consistently outperforms existing approaches in all but one of the dimensions analyzed. However, optimizing for an individual policy goal is often accompanied by unintended consequences as, e.g., focusing on efficient parking supply management alone leads to a high degree of inequity. Finally, we present a catalog of guidelines for practitioners and urban policymakers in Section 7.

2. Related work

Since the advent of paid parking, marked by the introduction of the first parking meter in Oklahoma City in 1935, priced parking has become a common sight in most centers of larger agglomerations around the world (Mingardo et al., 2015). Correspondingly, reflecting the rising scientific interest in its design and impact, various studies have proposed models of urban parking, in general, and priced parking, in particular (Inci, 2015). Following the classification proposed by Levy et al. (2013), this section serves to introduce a selection of these works before analyzing potential research gaps that this study attempts to address.

2.1. Spatially implicit and aggregate models of parking

The first class of modeling attempts has been made mainly from an economic perspective of parking with spatially implicit models that view the process from an aggregate perspective instead of modeling individual agents. For instance, Arnott et al. (1991) model traffic in a simple bottleneck model and find that location-dependent parking fees alleviate congestion caused by morning rush-hour traffic. In a similar vein, Shoup (2011) develops a model considering a series of variables, such as fees for on- and off-street parking,

[1] All code required for reproducing this study is available at https://github.com/JakobKappenberger/parking-fairness-paper.

parking duration, time spent searching for parking, and costs due to cruising, and demonstrates that even unpriced parking incurs societal costs, e.g., due to congestion stemming from cruising. Moreover, according to Shoup's model, cruising for parking is often a rational strategy when curbside parking is cheaper than adjacent off-street parking.

More recently, Hilvert et al. (2012) estimate a logit model of parking choice based on survey data and illustrate it using a simple case study. Macea et al. (2023), Yan et al. (2023), Weis et al. (2011) as well as Chaniotakis and Pel (2015) follow similar approaches of first conducting surveys and examining the importance of different factors, such as parking attitude, for the utility derived from a given parking alternative by estimating logit models. Jakob et al. (2020) evaluate a dynamic pricing scheme for parking fees in a macroscopic parking simulation in a test scenario representing the city of Zurich, Switzerland, to optimize the financial revenue generated for the municipality. van Nieuwkoop et al. (2016) propose a traffic flow model incorporating a search model for paid parking and investigate the potential efficiency gains of different pricing strategies for heterogeneous agents. Overall, since the macroscopic perspective taken for this class of models necessitates strong assumptions about the behavior of individuals and cannot perfectly account for the partial and stochastic characteristics of parking behavior, the validity of conclusions drawn based on these models can be somewhat limited (Levy et al., 2013).

2.2. Spatially explicit and individual models of parking

Consequently, beginning in the late 1990s, a growing number of studies presenting models that treat space explicitly and simulate drivers' behavior in a more detailed fashion were published. Pioneering this strand of work, Thompson and Richardson (1998) propose a parking search and choice model incorporating a utility function that drivers rely on for their decision-making as they navigate a simple street network one by one. They show that the experience gained during the search process does not significantly improve drivers' outcomes due to the high volatility of occupancy rates.

Building upon this, Benenson et al. (2008) introduce *PARKAGENT*, a spatially-explicit ABM modeled after Tel Aviv, Israel, and deploy the simulation to study the effects of additional parking supply in residential areas. Levy et al. (2013) propose both a second iteration of the *PARKAGENT* model (mainly with performance improvements) and *PARKANALYST*, a simpler analytical model that does not consider space explicitly, thus veering closer to the first category of parking models. Similarly, Dieussaert et al. (2009) present *SUSTAPARK* relying on the principles of cellular automata to simulate parking choices based on different utility functions for curbside and off-street parking. Waraich and Axhausen (2012) extend the traffic simulation framework MATSim to incorporate a parking choice algorithm based on a simple utility function. Similarly, Horni et al. (2013) create a cruising-for-parking component for MATSim following a cellular automaton approach. Rodríguez et al. (2022) improve upon these approaches by designing the empirically calibrated parking choice model *DYNAPARK* that can directly compare different types of parking. In Rodríguez et al. (2023), the authors extend their approach to include random coefficients that vary from one individual to another into their utility function. Tchervenkov (2022) implements a similar approach for MATSim, capturing the influence of different attributes of parking spaces not only on parking search but also on mode choice. Maxner et al. (2023) develop a parking simulation of Seattle in the simulation framework VISSIM to test different allocation configurations of curbside parking spaces with regard to their impact on various metrics, such as traffic flow or occupancy.

When analyzing these approaches, it is crucial to note that they all contribute essential advancements to the literature. Nonetheless, there are some interesting observations to make. While many of the mentioned studies model the parking behavior of individual drivers with some detail, the data used to calibrate these behavioral models is either not elaborated on (e.g., Waraich and Axhausen, 2012), older and might therefore not represent current behavior (e.g., *SUSTAPARK* relies on data collected in the 1980s), or based on rather small samples (e.g., *DYNAPARK* is based on 576 observations). Moreover, even though most of the studies examined include different types of parking, they typically model parking choice as a sequence of choices where drivers first decide on their preferred type of parking before choosing an individual alternative. However, it appears quite conceivable that drivers would prefer specific alternatives of a given parking type (e.g., a particular car park) over an alternative from a different type (e.g., parking curbside in a busy street) even though they may favor the latter type in most scenarios, thus rendering a joint comparison of all parking options potentially more realistic. Additionally, most of the simulations only allow for the local computation of the expected utility of a given parking alternative at the destination, despite surveys showing that some respondents pursue the strategy of searching for parking en route to their destination (Polak and Axhausen, 1990).

Finally, and of particular relevance to this study, none of the examined parking models investigate the fairness of parking as these approaches mostly focus on traditional performance-related metrics, such as traffic flow due to cruising for parking, occupancy, or generated revenue. If an effort is made to optimize these metrics, simple approaches, such as conditional rules for dynamic pricing, are deployed.[2] However, given the intricacies of parking behavior, more complex and potentially better-performing methods such as ML might provide further insights into the (unintended) consequences of pricing strategies. Therefore, building upon this literature review, our goal was to create a simulation of urban parking addressing these research gaps by featuring a complex behavioral model of parking choice that is calibrated using a large and representative survey sample and allows for the simulation of different parking strategies. In doing so, we aim to evaluate different ML-based pricing schemes for parking against established baselines regarding their effects on fairness and commonly deployed performance measures that reflect different policy goals. The following sections 3, 4, and 5 present our approach in detail.

[2] The predecessor to the simulation presented in this study is an exception to this observation, as it is employed to test ML-based pricing schemes. However, the behavioral model governing the simulation is rather primitive compared to the studies presented (Kappenberger et al., 2022).

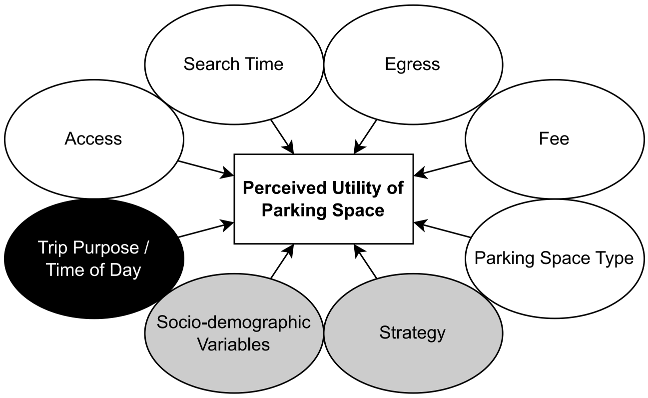

Fig. 2. Parking choice model. White circles designate parking space-specific factors, gray circles show individual-specific factors, and black circles describe trip-specific factors.

3. Modeling parking behavior

In order to create an empirically validated model of parking choice behavior, we designed and conducted a survey featuring a Discrete Choice Experiment (DCE). Beforehand, we formulated a theoretical model of parking choice, which is explained in Section 3.1. On this basis, we created the survey design (Section 3.2). The data gathered from the survey was then used to estimate a logit model of parking choice (Section 3.3) that is fed into the simulation model described in Section 4.

3.1. Defining a model of parking preferences

As it represents the most widely used behavioral paradigm for modeling transport-related choices, in general, Cascetta (2009) and parking behavior, in particular (e.g., Axhausen and Polak (1991), Waraich and Axhausen (2012), or Antolin et al. (2018)), the parking choice model deployed in this study is based on Random Utility Theory (RUT). RUT rests on the assumption that a given individual i acts rationally, choosing the alternative j with the highest perceived utility (or attractiveness) U^i_j out of a set of potential options $\boldsymbol{I}^i$. This utility depends on attributes of the alternative as well as the individual itself that can be expressed in the vector $\boldsymbol{X}^i{}_j$ and mapped via a utility function U^i to U^i_j. Since the perceived utility of individual decision-makers cannot be estimated with certainty, the random error term $\varepsilon^i{}_j$ is deployed to model the residual error due to variance not covered by $\boldsymbol{X}^i{}_j$ (Domencich and McFadden, 1975; Cascetta, 2009). Consequently, to construct $\boldsymbol{X}^i{}_j$, it is crucial to identify the attributes relevant to the specific choice situation, in this case, parking. Fig. 2 illustrates the factors incorporated in the model used for this study.

Following the work of Axhausen and Polak (1991), and others before them (Ellis et al., 1974), the basis of the model is formed by well-established attributes of the individual parking spaces, namely *access* (the time spent driving to the parking space), *search time* (the time spent searching for an unoccupied space at location), *egress* (the time spent walking to the final destination after parking), the *type of parking* (curbside versus off-street parking facilities) and, finally, the *fee* incurred by parking at the parking space in question. Beyond that, we also include factors that vary not by parking space but by individual or trip.

Similar to Waraich and Axhausen (2012), who incorporate drivers' incomes in relation to the fee of parking at a respective parking space into their utility function, our model allows for the inclusion of socio-demographic variables that may correlate with different types of parking behavior. In our case, based on previous findings in the literature (Brooke et al., 2014), these encompass *household income*, *age*, as well as *gender*. In a similar fashion, we consider trip-specific attributes, such as *trip purpose* or *time of the day*, as influences since they might impose different constraints on drivers' parking choice behavior (Brooke et al., 2014).

Moreover, the model differentiates between a series of general archetypes of parking search behavior that were identified by Polak and Axhausen (1990) and codified in the following strategies: (1) always using the same parking space, (2) having a private or reserved parking space, (3) starting to look for parking after arriving at the destination, (4) parking in the closest car park to the destination, or (5) taking the first unoccupied space while driving to the destination. We added an additional strategy (6) of using apps to identify and pay for parking spaces to account for the increasing popularity of app-based solutions for finding parking spaces in city centers. Furthermore, there is a category (7) for "other" strategies. These parking strategies are meant to account for variance in an individual's parking behavior unexplained by the attributes of the individual, the trip, or the parking space in question.[3]

Before proceeding, it is worth noting that the RUT framework has received a fair share of criticism over the years (see Hess et al. (2018) for a detailed discussion of these potential shortcomings). The criticism primarily focuses on the observation that, while RUT seemingly requires individuals to act perfectly rationally (from the view of the researcher in question), the empirical reality of

[3] Importantly, in contrast to e.g., Polak and Axhausen (1990), we only include legal parking in this first iteration of our approach.

Table 1
Mixed Logit Model of Parking Choice. Interaction effects that are not statistically significant are not included (see Table A.7 for the complete model). σ refers to the standard deviation of the coefficient mean if random coefficients are estimated.

	Coefficient	σ
Access (min)	−0.04 (0.02)*	−0.04 (0.04)
Search Time (min)	−0.05 (0.04)	0.14 (0.04)***
Egress (min)	−0.24 (0.04)***	0.20 (0.03)***
Space Type Car Park	−0.05 (0.18)	0.80 (0.19)***
Fee (€)	−1.23 (0.26)***	0.84 (0.08)***
Age	−0.00 (0.00)	0.01 (0.00)**
Gender Female	0.16 (0.10)*	0.34 (0.33)
Strategy En Route * Egress	0.10 (0.03)***	–
Strategy Car Park * Space Type Car Park	0.87 (0.16)***	–
Strategy Car Park * Fee	0.52 (0.08)***	–
Strategy Other * Search Time	−0.09 (0.04)**	–
Strategy Other * Egress	0.12 (0.04)***	–
Time Afternoon * Space Type Car Park	0.25 (0.15)*	–
Purpose Doctor * Search Time	−0.07 (0.04)*	–
Purpose Doctor * Egress	−0.09 (0.03)***	–
Purpose Doctor * Fee	0.68 (0.09)***	–
Purpose Acquaintance * Access	−0.08 (0.03)***	–
Purpose Acquaintance * Egress	−0.09 (0.03)**	–
Purpose Acquaintance * Space Type Car Park	0.35 (0.18)*	–
Purpose Acquaintance * Fee	0.24 (0.09)***	–
Purpose Shopping * Egress	−0.15 (0.03)***	–
Purpose Shopping * Space Type Car Park	0.31 (0.17)*	–
Purpose Shopping * Fee	0.52 (0.09)***	–
Household Income 2 * Fee	−1.01 (0.48)**	–
Household Income 3 * Fee	−0.74 (0.27)***	–
Household Income 4 * Fee	−0.86 (0.25)***	–
Household Income 5 * Fee	−0.75 (0.25)***	–
Household Income 6 * Fee	−0.74 (0.25)***	–
Household Income 7 * Fee	−0.67 (0.25)***	–
R^2_{McFadden}	0.37	
Log Likelihood	−2764.19	
Num. obs.	6301	

$^{**}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

human behavior appears to be at odds with this premise. We still opted to utilize RUT but attempt to account for such inconsistencies by estimating random coefficients and including error terms in the behavioral model when deploying it in our simulation, conceding that no behavioral model represents a one-size-fits-all solution (see Section 3.3 as well as Section 4.2 for details).

3.2. *Survey: Experimental design*

To validate the parking behavior model and quantify the effects of the individual factors, we developed a DCE as part of an online survey where respondents were asked to select one of two parking alternatives. DCEs are a commonly deployed methodology to study choice behavior by repeatedly offering respondents multiple alternatives with varying attributes to relate the choices made to these factors and attributes of the person (Friedel et al., 2022). For this purpose, they are also commonly used for modeling parking choices (e.g., Axhausen and Polak (1991), Hilvert et al. (2012), Rodríguez et al. (2022)). We provide a detailed description of the design of our DCE in Appendix A.1.

3.3. *Logit model*

The field phase of the survey lasted from May 3 to May 23, 2023. The survey was run online and completed by 2,021 respondents recruited via a probability-based online panel designed to represent the German (online) adult population. We applied crossed quotas for gender and age groups (see distribution in Table A.4).[4] After removing "speeders" (i.e., respondents that were quicker than 60 percent of the median duration of survey completion (11 m 15s)[5] (Roßmann, 2010)) and excluding non-answers for the variables to be investigated, as well as respondents without a driver's license, 1,578 respondents providing data on 6,301 choice tasks remained. Since some of the parking strategies detailed above were scarcely selected (e.g., only 37 respondents chose strategy 6), we merged strategies 1, 2, and 6 into strategy 7 ("other").

[4] Table A.5 shows the distributions of the remaining variables included in our model.

[5] We acknowledge that survey duration can only serve as a proxy of response burden since the burden perceived by respondents is both subjective and multi-dimensional (Yan et al., 2020).

For data analysis, we used R with the package *mlogit* to estimate a mixed logit model of parking choice (Croissant, 2020). Mixed logit models are a variety of logit models that can account for the heterogeneity within the population by allowing the model parameters to vary from one individual to another. Moreover, these models can incorporate the nested structure of our data as we have multiple observations (i.e., choice tasks) per individual. Instead of estimating an individual coefficient per person, mixed logit models estimate the distributions of the coefficients to obtain random coefficients (Train, 2009). Since we assume a normal distribution of the coefficients, this equates to estimating the standard deviations σ of the coefficients in addition to the coefficient mean. To obtain a somewhat parsimonious model, we estimate random coefficients only for the main effects in our model.[6]

Regarding modeling the different strategies, we opted against estimating separated models per strategy as this would have significantly reduced the sample size per model. Instead, as it appears quite likely that different strategies may correlate with individual attributes, we included interaction effects between the strategies and the individual attributes, omitting their main effects. Moreover, we hypothesize that household income affects the perceived importance of parking fees since higher-income households most likely place less importance on fees than their lower-income counterparts. Consequently, we include the interaction effects between the parking fee of a parking space and the categorical variable household income (coded in seven income groups, ranging from up to €520 for group 1 to €5000 and beyond for group 7).

Table 1 displays the mixed logit model of parking choice. Due to space limitations, the table only contains the main effects as well as the statistically significant interaction effects. For the complete results, see Table A.7 in Appendix A.2. Overall, while not comparable between different models, the R^2_{McFadden} value of 0.37 indicates a very good model fit (McFadden, 1979).

The coefficients of the parking space attributes behave quite similarly to those reported in the literature (this holds in particular for the basic model containing only the main effects shown in Table A.6 in Appendix A.2).[7] While search time and space type are no longer statistically significant influences after accounting for interaction effects, access and particularly egress are relevant to respondents' decision-making. In particular, as already reported by Axhausen and Polak (1991), respondents place much more importance on avoiding walking as opposed to the time spent in the car on their way to their destination. Predictably, the fee incurred for parking at a particular parking space is negatively associated with the choice of a parking alternative and represents the clearly strongest parking space-specific predictor for parking preference.

Proceeding to the different parking strategies, where strategy 3 ("Close to Goal") represents the reference category, respondents identifying with strategy 5 ("Èn Route") are more willing to accept longer walks to their destination, as evidenced by the positive interaction effect with egress. This appears reasonable since finding parking earlier en route to the destination should correspond to higher egress when compared to starting to look once arrived. As expected, respondents with strategy 4 ("Car Park") indicated a preference for parking in car parks and a willingness to pay more for parking when compared to drivers following the other strategies. Overall, these findings indicate that the approach of utilizing interaction effects to model the different approaches to parking succeeded in identifying meaningful differences between them. While there are significant interactions between the "other" strategy and search time as well as egress, it is difficult to interpret them as they cannot be attributed to any individual strategy.

There is only one statistically significant effect to report involving the time of day of a given trip to the city center, as respondents were more keen on parking in car parks in the afternoon. Conversely, for the trip purposes, there are statistically significant differences between the listed purposes and the reference category "Work/Education". In particular, respondents are willing to pay more for parking for any other purpose and emphasize avoiding walking longer distances. Moreover, for the purposes "Acquaintance" and "Shopping"', they also showed a stronger preference for parking in car parks.

Finally, the interaction effects between the different household income groups in the survey and the parking fee confirm our previous hypothesis. Apart from the fourth income group, which nonetheless also has a negative interaction effect with parking fees, the importance of the parking fee decreases linearly with rising household income. However, given that the coefficients of the income groups are negative compared to the reference category ("1"), the lowest income group in the survey, respondents of said group actually indicated less trepidation to choose higher-priced alternatives in the survey than those belonging to higher income strata. This may be due to the small number of 13 respondents indicating a household income in this range, potentially resulting in statistical noise.

Having formulated our parking behavior model and estimated its coefficients, it remains to introduce the parking simulation incorporating said behavioral model.

4. Simulation

The visual interface of the ABM for parking that forms the center of this study is shown in Fig. 3. The simulation presented here is based on the model described in Kappenberger et al. (2022), extensively modifying and extending the behavioral rules as well as a large number of other subroutines governing the previous version. The simulation is implemented in *NetLogo* (Wilensky, 1999) and modeled after the city of Mannheim, Germany. The municipality provided us with empirical data on traffic volume and parking demand. We used this data as well as the results of the survey on parking behavior featured in the previous section to calibrate our model. In the following, we will present the model's environment (Section 4.1), its agents (Section 4.2), and the behavioral rules governing their actions (Section 4.3).[8]

6 For the categorical variable *gender*, the value "diverse" had too small a sample size to allow inclusion into our models.

7 While the coefficient estimated for search time in relation to access and egress is in line with recent studies (e.g., Rodríguez et al. (2022)), it is relatively low compared to previous contributions, such as Axhausen and Polak (1991). This may be caused by potential difficulties respondents may have had differentiating between access and search.

8 Appendix B.1 explains the calibration process of the simulation.

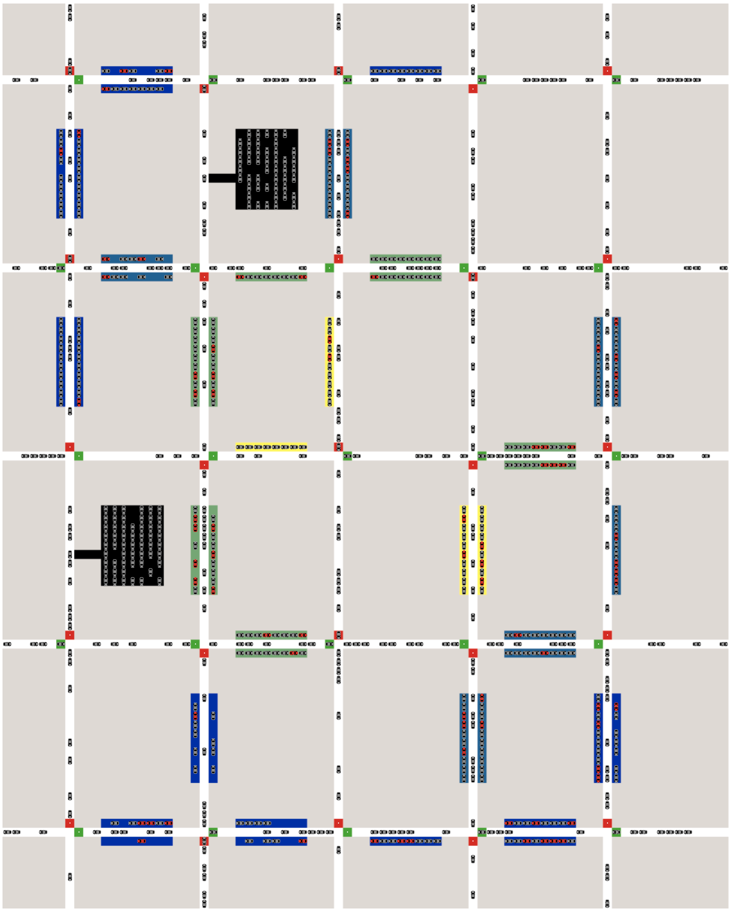

Fig. 3. Visual interface of the simulation for parking.

4.1. Environment

The model's environment consists of the traffic grid (see Fig. 3), which resembles the street layout of the city center in Mannheim. Situated at the curbside of the roads, the yellow, green, teal, and blue patches designate four different Controlled Parking Zones (CPZs) that are grouped according to their distance to the center of the map. The larger, black squares represent off-street parking facilities. The positions of these parking alternatives are fixed for every model run. Each of these runs lasts for 21,600 time steps, corresponding to a simulated time frame of 12 h from 8:00 a.m. to 8:00 p.m.

4.2. Agents

The agents of our model are individual cars moving across the grid. The initial number of agents varies around a mean value (*num-cars-mean*) with each model run in accordance with the mean variation found in traffic counts of our model city (around ± 10 percent). Moreover, the relative traffic volume in the simulation changes over the course of the day in accordance with the distribution shown in Fig. B.10 in Appendix B, similarly derived from traffic counts in Mannheim. The drivers in the model have a series of attributes that are based on a synthetic data set created based on the data gathered from the survey described in Section 3.3. The data set was created using the *DataSynthesizer* library to ensure the preservation of the empirical distributions of the relevant attributes (and their correlations) while protecting the privacy of respondents (Ping et al., 2017). Among these are age, gender, household income, and parking strategy. Moreover, following the OECD (2019), we added a simplified income classification to group agents into three income groups. Complementing the seven income groups in our survey data, incomes between 75% and 200% of the median income $\tilde{x}$ (calibrated based on income data in Germany) are assigned the "middle-income" label. Deviations above or below this mark are designated "high-" or "low-incomes", respectively. Lastly, all agents are assigned a parking duration, which, in accordance with the findings of Jakob and Menendez (2021), is modeled as following a gamma distribution. The speed of the vehicles is calculated as patches traveled per tick in our Netlogo model. The average speed of all agents is then computed using only the speeds of the cars that are not parked. Due to technical reasons, the maximum speed is approximately 0.9 patches per tick,

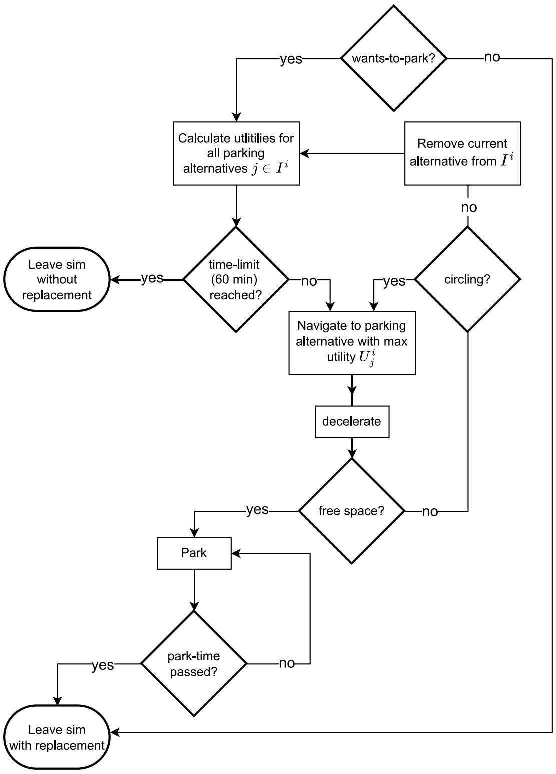

Fig. 4. Flow diagram visualizing the behavioral rules of the simulation.

as higher values lead to rounding errors in NetLogo when determining the current patch of the respective vehicle. We then further assume that the maximum speed in the model corresponds to the maximum speed in the city center of Mannheim (30 km/h).

4.3. Behavioral rules

Fig. 4 shows the behavioral logic governing the actions of the individual agents. All cars navigate the grid with previously assigned destinations.[9] For those simply traversing, this destination is one of the exit points of the street network. For cars seeking to park, destinations are determined with probabilities inversely proportional to their distance to the center of the grid, accounting for the higher popularity of the center. The share of traffic trying to park is determined by the distribution formed by the parameter *demand-curve-intercept* and data on parking demand in Mannheim (based on occupancy data provided by the municipality and shown in Fig. B.10 in Appendix B).

Generally, for all agents attempting to park, we differentiate between their perceived utilities and their achieved outcomes. They calculate the perceived utilities U_j^i of all parking alternatives $j \in \boldsymbol{I}^i$ available to them using their individual utility function U^i. U^i derives the utility of a given parking alternative by computing a weighted sum of the alternative's attributes and the interactions between Agent i's parking strategy, household income, and trip purpose with these attributes, collected in vector ${\boldsymbol{X}^i}_j$. The individual utility weights $\boldsymbol{W}^i$ assigned to every car after spawning are based on the coefficients estimated by the logit model described in Section 3.3. For the random coefficients, weights are drawn for every agent according to the coefficient mean and standard deviation. Only significant coefficients are included in $\boldsymbol{W}^i$ (i.e., all coefficients listed in Table 1), resulting in the following definition of U_j^i for any parking alternative $j \in \boldsymbol{I}^i$:

$$U_j^i = U^i({\boldsymbol{X}^i}_j) = {\boldsymbol{W}^i}^\top \cdot {\boldsymbol{X}^i}_j + {\varepsilon^i}_j = (\sum_{i=1}^{k} {\boldsymbol{W}^i}_k {\boldsymbol{X}^i}_{jk}) + {\varepsilon^i}_j,$$

where k is the length of both $\boldsymbol{W}^i$ and ${\boldsymbol{X}^i}_j$ and ${\varepsilon^i}_j$ constitutes the residual error term that is, in accordance with the general assumptions of multinomial logit models, drawn from a Gumbel distribution and modeled as an independent and identically distributed random variable to account for unobserved heterogeneity in parking behavior (Train, 2009).

[9] Their distribution can be inspected in Fig. B.9 in Appendix B.

While parameters such as *egress* or *fee* can be simply queried by the agents, their expected *search-time* has to be estimated. Therefore, when computing U_j^i, agents rely on information about the amount of traffic on their route, similar to what navigation apps are capable of producing. Depending on this value, a penalty is added to the estimated *access* value since it is assumed that more traffic corresponds to a higher likelihood of longer search times. The actual search time is then recorded as the time spent searching locally after the street of the first parking alternative chosen was reached.

After completing these computations, the agents navigate to the parking alternative, offering them their maximum utility U_j^i via the shortest route available. Once a car has entered the street with its parking location of choice, it decelerates and parks on any unoccupied parking space, as all parking alternatives on a given block are treated as equal. The computed utility of this parking space, updated with the recorded values for search time and access, is then documented as the outcome O^i achieved by the agent.

Should all parking spaces turn out to be occupied, drivers will recompute U_j^i of all potential alternatives except the one just visited since it is removed from the choice set I^i, as attributes such as egress, access, and fee depend on the location of the agent or may change on their own. An exception to this behavior present agents that engage in circling. These agents will circle the block once per selected parking alternative, hoping a space will become available in the meantime.[10] All agents will then once more drive towards the parking space that offers the maximum U_j^i (although for circling drivers, this parking space remains the same). This procedure is repeated until an hour of modeled time has passed since it is assumed that drivers will then resort to finding parking outside the modeled area. Moreover, these agents are assigned the minimum outcome of any agent currently parking in the simulation based on the assumption that their outcome O^i is at least as unfavorable as the worst-off parking in the targeted area.

Upon completion of their parking time, cars leave the CPZ and navigate towards the edge of the grid, where they are replaced with newly set up cars. In contrast, cars unable to find parking are not replaced once they leave the map, leading to a decrease of *num-cars-mean*. This assumes that the modeled drivers permanently change their strategy to refrain from trying to park in the city center, preserving the change to the social distribution in the model that this behavior introduces.

5. Operationalizing the fairness of parking fees

Optimizing parking pricing only with respect to single aggregate policy goals, such as reduced traffic and emissions, may lead to important unintended consequences for unfairness and social inequality. However, as argued in Section 2, fairness considerations have not been heavily featured in the literature on parking and parking simulations. In the present paper, we therefore make use of ML not only as it promises more powerful parking space use optimization. Given that ML is also particularly prone to generating or reinforcing biases, we also draw specifically on "fair machine learning" (Barocas et al., 2019) research that aims to ensure fair outcomes, e.g., by considering how prediction outcomes vary for different subgroups of the population. To this end, a series of metrics have been proposed to quantify bias (Mehrabi et al., 2022) that can also be generalized for our use case and will thus serve as a starting point for these considerations.

Many of these metrics are variants of the principle of *fairness by unawareness*, which requires that forbidden attributes – often membership in a particular, for instance, ethnic or economic subgroup – are not explicitly used in the decision-making process (Grgic-Hlaca et al., 2016). This idea that people from different groups should be treated in the same way thus implements the concept of *equality*.

However, we aim to operationalize fairness by relying on the construct of *equity*. In contrast to the concept of equality that suggests treating all citizens in the same way (e.g., by providing the same level of services to all citizens), equity also considers individual circumstances. It suggests that citizens should be treated in such a way that the individual outcomes are as similar as possible, for instance, by providing a higher service level to citizens with more needs (Deutsch, 1975). Therefore, equity focuses on the outcome side of a policy instead of the treatment side (which is central to equality measures). Equity-based metrics are therefore suited for our approach, which is interested in the outcomes that agents from different income groups achieve with a given pricing policy.

Any given policy is unlikely to achieve complete equity in realistic scenarios with complex outcome structures. Instead of regarding equity as a binary property, it is, therefore, more suitable to measure the degree to which a policy achieves equity. For this purpose, we define the degree of individual inequity of a policy (i.e., in our case, a pricing policy for parking fees) as the divergence between the distribution of outcomes for all persons affected and the discrete uniform distribution $\mathcal{U}$, which would correspond to equal outcomes for all. All results reported in this paper have been computed using Jensen-Shannon Divergence (JSD) as a divergence measure, which we normalize to a maximum value of 1 using the total variation distance as it serves as an upper bound for the JSD. The JSD is commonly used in statistics and measures the similarity between two probability distributions (or, in our case, outcome distributions).[11] Lower values indicate more similarity. For better readability of the results, we take the square root of the JSD, which effectively gives the Jensen–Shannon distance.

In real-world applications, however, it is often not practicable to determine individual inequity simply because it is infeasible to identify the strategies and preferences of each individual involved. Therefore, we focus on *inter-group inequity* that we define as the inequity between the demographic groups under consideration. For the sake of simplicity, we use the average outcome of group members as a basis for the measure. In the present case, the compared groups are the three income groups as defined above. While many further group attributes may be affected by pricing policies, income is arguably one of the most directly relevant variables

[10] In accordance with empirical work, we assign ten percent of drivers the attribute *circling?* triggering this behavior (Montini et al., 2012).

[11] See Menéndez et al. (1997) for a detailed definition and discussion of the JSD.

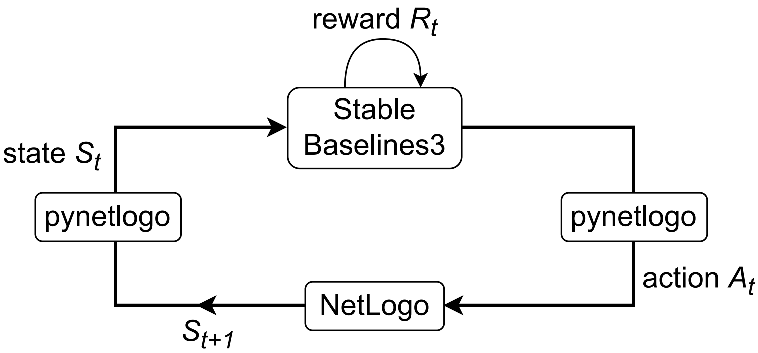

Fig. 5. Interaction between Stable Baselines3 *(RL Agent)* and NetLogo *(RL Environment)*.

for pricing policies. This approach is theoretically and practically useful. Theoretically, it allows us to learn about mechanisms that lead to unequal outcomes of different pricing policies between income groups. This is also of practical relevance for policymakers to improve parking policies so as not to disadvantage specific social groups. The resulting measure is defined as follows:

Definition 5.1 (*Inter-Group Inequity*). Let $f_{[p_k]} = avg\{O^k | p_k \in [p_j]\}$ be the average outcome of persons in equivalence class $[p_j]$ and let $F(\sim) = \left[f_{g_1}, \dots, f_{g_{|N/\sim|}}\right]$ be the average outcomes of the n equivalence classes, then inter-group inequity with respect to equivalence relation $\sim$ is defined as

$$inequity(\sim) = JSD\left(\frac{F(\sim)}{\sum_{j=1}^{|N/\sim|} f_{g_j}} \parallel \mathcal{U}(n)\right)$$

Thus, we measure the similarity between the observed outcome distribution of the different equivalence classes and the discrete uniform distribution $\mathcal{U}$, which would give each group equal outcomes. Obviously, individual inequity is a special case of inter-group inequity, where each individual forms their own group. As argued above, inter-group inequity is the most suitable measure for analyzing inequity in complex real-world situations.

6. Experiments: Results and discussion

The following section first serves to introduce our experimental setup with a particular focus on the ML-based pricing system deployed before describing and discussing the results of our experiments.

6.1. Experimental setup

To introduce ML-based pricing for parking to our NetLogo ABM, we relied on Reinforcement Learning (RL). RL represents the third paradigm of ML next to supervised and unsupervised ML. The concept can be described as "learning what to do—how to map situations to actions—so as to maximize a numerical reward signal" (Sutton and Barto, 2018, p. 1). As the learning agent is not given the set of actions that benefit its goal of maximizing its reward, it has to proceed by trying different actions and studying their impact on the environment. In doing so, it must strike a balance between *exploiting* the knowledge it has already obtained and *exploring* new strategies to test them for potential gains. Due to this dynamic type of learning, RL is well suited to be applied in a pricing system for parking to react dynamically to shifting circumstances. Before proceeding, it is important to note that there is significant overlap between the terms used for RL and ABMs. Thus, when we refer to one of the components constituting a RL system, we will prefix said term with "RL".

Fig. 5 illustrates the general workflow of the RL-based pricing scheme. The *pynetlogo* library is deployed to communicate with NetLogo's API to control the simulation from within a Python session (Jaxa-Rozen and Kwakkel, 2018). As RL framework, we used *Stable Baselines3* (Raffin et al., 2021).

We designed a custom RL environment as a wrapper for the ABM and a communication interface for the RL agent. The custom environment receives the current state of the traffic simulation, computes the appropriate reward, and sends the actions determined by the agent back via pynetlogo. A time step t occurs in intervals of 900 ticks in the model, allowing the RL agent to adjust prices every 30 min simulated. At every time step t, the RL agent queries the current state S_t from the RL environment, simulated by NetLogo. Table C.9 in Appendix C displays the information contained in S_t.

Based on these numerical state representations, reward R_t is calculated according to the reward function supplied to the agent (an overview of reward functions used in this paper is provided below). Completing time step t, the next actions A_t are transmitted via pynetlogo to the ABM. Individual actions can originate from an integer range from 0 to including 20. The action values are then divided by 2 to freely set fees for a CPZ between €0 and €10 in €0.50 intervals. This strategy has the advantage of enabling immediate reactions to occupancy changes and, conceptually, an easier process of credit assignment for the agent, allowing it to better anticipate drivers' reactions to its pricing strategy. In scenarios where we allow group-specific pricing, the RL agent determines distinct fees per income group for every CPZ. After the fee changes have been implemented in NetLogo and the simulation has continued to run for another half hour of simulated time, the environment sends the new state S_{t+1}. This process repeats until the episode is complete, prices have been adjusted 24 times, and the terminal state has been reached. Thereafter, the environment in NetLogo is reset, and the next episode begins. In our experiments, we utilized the following reward functions to compute R_t:

$r_{occupancy}$. Existing dynamic pricing policies like the SFpark system in San Francisco are designed to reach a certain level of occupation of on-street parking. A typical goal is a utilized capacity of 80%, which corresponds to a good level of use while still leaving space for newly arriving cars (Shoup, 2011). We use an objective function that rewards occupancy levels between 75% and 90% and punishes the agent with increasing severity for moving outside this range.

r_{equity}. To examine how the learner behaves if it is tasked with specifically maximizing equity, we use a reward function that encourages the agent to minimize the inter-group inequity between the different income groups in the simulation. As initial experiments have indicated a performance improvement, the function is based on the outcome difference between the best- and worst-off income groups rather than directly optimizing for inter-group inequity.

r_{equity} *with group-specific pricing.* As it is likely that managing demand with dynamic prices equal across all income cohorts will disproportionately impact those of low income (at least if fees are not consistently kept at €0), we designed an experimental setting in which the agent is also tasked with minimizing inter-group inequity. However, it is now able to set group-specific prices at the different CPZs. In theory, this should enable the RL agent to achieve a higher degree of equity as it is able to tailor its actions specifically to the social groups analyzed. While this approach most likely sacrifices some empirical validity, as it would be challenging to implement in practice, proxies, such as car size, have already been proposed as a basis for pricing parking spaces.

r_{speed}. This reward function maximizes the average speed of the non-parking cars driving in the grid at time step t. It is meant to allow the agent to pursue the goal of increasing mobility in the modeled municipality. Higher average travel speed corresponds to more ease of movement for the cars in the simulation. Conversely, lower values for the metric will most likely be accompanied by more traffic volume and congestion, lowering overall mobility. Moreover, increased traffic flow also correlates with lower emission levels, as congestion is accompanied by higher pollutant and noise emissions (Zhang et al., 2011), negatively impacting the quality of life of city dwellers and visitors alike.[12]

$r_{composite}$ *with group-specific pricing.* Finally, we combine $r_{occupancy}$ with the group-specific pricing version of r_{equity}. The former optimizes for the variable that the system has the most direct influence on and represents one of the premier goals pursued by municipalities in this regard. The latter accounts for the fairness of parking. With this reward function, we aim to investigate how well the learner can balance different, potentially conflicting policy goals.

We deployed the Proximal Policy Optimization (PPO) algorithm, one of the most commonly used RL algorithms, to maximize these reward functions (Schulman et al., 2017). Furthermore, we conducted Bayesian hyperparameter tuning for all reward functions introduced by testing different parameter vectors over 16 iterations of training for 20,000 episodes for each function. For reproducibility, our final hyperparameter configurations per reward function are listed in Appendix C, Table C.10. Training was then conducted for 50,000 episodes using the smaller version of our ABM,[13] followed by a fine-tuning period of 5,000 episodes of training on the evaluation model to ease the transfer of the pricing policy learned.[14]

Moreover, as appropriate benchmarks for our ML-based pricing system, we deploy two baselines:

Static baseline. To mirror the predominant pricing strategies in many municipalities today, we implemented a *static baseline* pricing, consisting of fees of €3.5 per hour for all CPZs on the map. This corresponds to the pricing scheme active in our model municipality.

Dynamic baseline. To cover more modern pricing strategies, we deploy a simple dynamic pricing scheme as a *dynamic baseline*, which we model as a more demand-responsive version of *SFpark*, the system active in San Francisco: A CPZ's fee is raised by €0.25 if more than 90% is occupied. Conversely, prices are lowered by €0.25 when the individual utilized capacity falls under 75% and by €0.50 if it amounts to less than 30%. Similar to the ML-based approach, adjustments are made every 30 min.

Both the different ML learners optimizing the varying reward functions as well as the baselines were then run for 100 iterations on the evaluation model to achieve robust results for analysis.

6.2. Results

Before examining the results achieved by the individual pricing schemes, we first evaluate the performance of the behavioral model of parking underlying the ABM deployed. In particular, Figs. 6(a) and 6(b) illustrate the degree to which the different parking strategies correspond to distinct behavioral patterns in the simulation. As indicated by the logit model in Section 3.3, on average, agents following the strategy "En Route" park in parking spaces that are further away from their final destination. This does not

[12] This also depends on traffic volume and only holds up to certain speed levels (over 100 km/h and thus not relevant here), as the relationship between car speed and many traffic-related emissions is u-shaped (Int Panis et al., 2006).

[13] See Appendix B.1 for detailed information on the different simulation sizes used as well as their calibration.

[14] Running 64 simulations in parallel on an AMD Milan EPYC 7513 processor, training took approximately 19 h for $r_{occupancy}$.

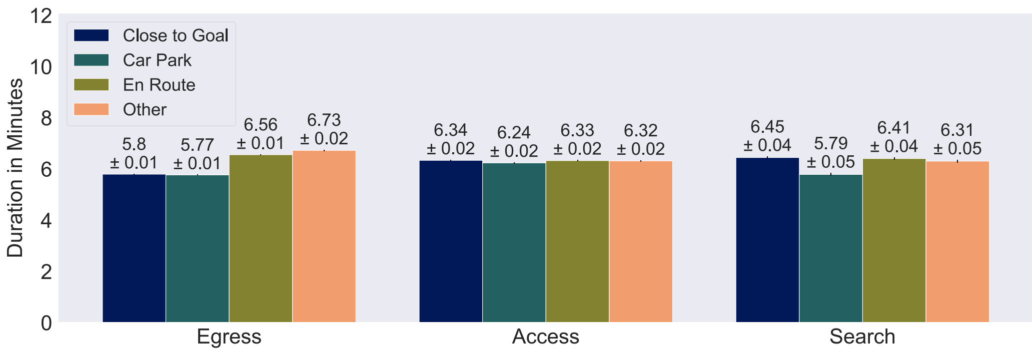

(a)

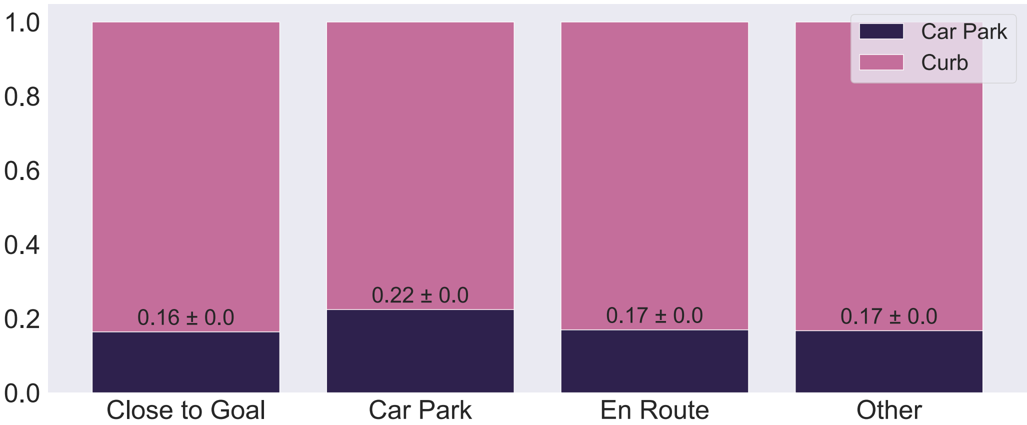

(b)

Fig. 6. (a) compares the attributes of the parking spaces for drivers following the different parking strategies. (b) illustrates the share of drivers parking in off-street parking facilities for all parking strategies. (both aggregated over all evaluation episodes of all strategies examined, error bars indicate 95% confidence intervals).

result in a lower average value of access, however, as the confidence intervals of the values of the strategies "En Route" and "Close to Goal" overlap, thus indicating that drivers belonging to the former do not spend significantly less time in their cars on average. Expectedly, the third distinct strategy, "Car Park", results in its followers parking significantly more often in car parks compared to the other strategies. Overall, these observations show that the behavioral model successfully nudges the agents toward the expected parking choices.

We now turn to the overall performances of the different pricing strategies deployed. Fig. 7 visualizes their performances relative to one another, and Table 2 quantifies them on an absolute scale (see Appendix D for more detailed results). First, it is worth noting that, in all but one dimension, the baseline scores were improved by a ML-based system. The *static baseline* confirms the shortcomings of static pricing schemes mentioned in the literature as it exhibits a relatively poor performance regarding optimizing the utilized capacity of the CPZs as well as lowering overall traffic volume. Nonetheless, while the constant pricing of the baseline appears ill-equipped to manage parking demand, it leads to the highest average speed and, perhaps less surprisingly, a relatively low level of inequity among the different income groups, only behind learners that optimize for this dimension. Conversely, the *dynamic baseline* performs better in keeping the occupancy of the curbside parking supply in the desired range but shows worse results across all other categories. In particular, the high fees resulting from the dynamic strategy lead to the third-highest degree of inequity among all pricing systems, thus demonstrating that the improvement in parking supply management offered by such systems may come at a cost. Regarding the surprising finding of the constant fees leading to better traffic flow, we hypothesize that this is due to the fixed prices allowing for a more even distribution of agents in the model because they do not have to react to constant price shifts, rendering different CPZs the most popular in the simulation and leading to a concentration of traffic in and around the respective zone. In this respect, it is essential to note that while cruising traffic is slower than through traffic (as is the case for all parking strategies), our experiments do not cover the long-term consequences of a given parking strategy on the relationship between these two kinds of traffic as we only simulate individual days. Thus, it is well possible that if implementing one of the examined alternatives, we might observe more favorable traffic flow compared to the *static baseline* once drivers have fully adapted their behavior to a given strategy by, e.g., switching their preferred mode of transport in light of higher fees.

Table 2

Results achieved by the different pricing strategies over 100 evaluation episodes. "Occup." is the timeshare that CPZs are held in the desired occupancy range. "Traffic Flow" is the average normalized speed, and "Traffic Volume" corresponds to the average amount of cars relative to the starting value per episode. In the "Outcome Averages" category, "Overall" refers to the average outcome among all agents, while "Low", "Middle", and "High" mark the average outcomes in the respective income groups.

		Traffic		Outcome Averages				
Strategy	Occup.	Flow	Volume	Overall	Low	Middle	High	Inequity
Static Baseline	0.38	**0.49**	0.96	−2.88	−4.61	−2.66	−1.33	0.301
Dynamic Baseline	0.40	0.49	0.96	−2.89	−4.74	−2.65	−1.19	0.329
$r_{occupancy}$	**0.56**	0.47	0.95	−2.8	−4.92	−2.52	**−0.82**	0.407
r_{equity}	0.28	0.47	0.95	**−1.52**	**−2.06**	**−1.44**	−1.08	0.166
r_{equity} (group)	0.34	0.45	**0.95**	−1.92	−2.11	−1.99	−1.29	**0.13**
r_{speed}	0.39	0.49	0.96	−2.77	−4.69	−2.5	−1.15	0.337
$r_{composite}$	0.55	0.46	0.95	−2.3	−2.48	−2.52	−0.96	0.107

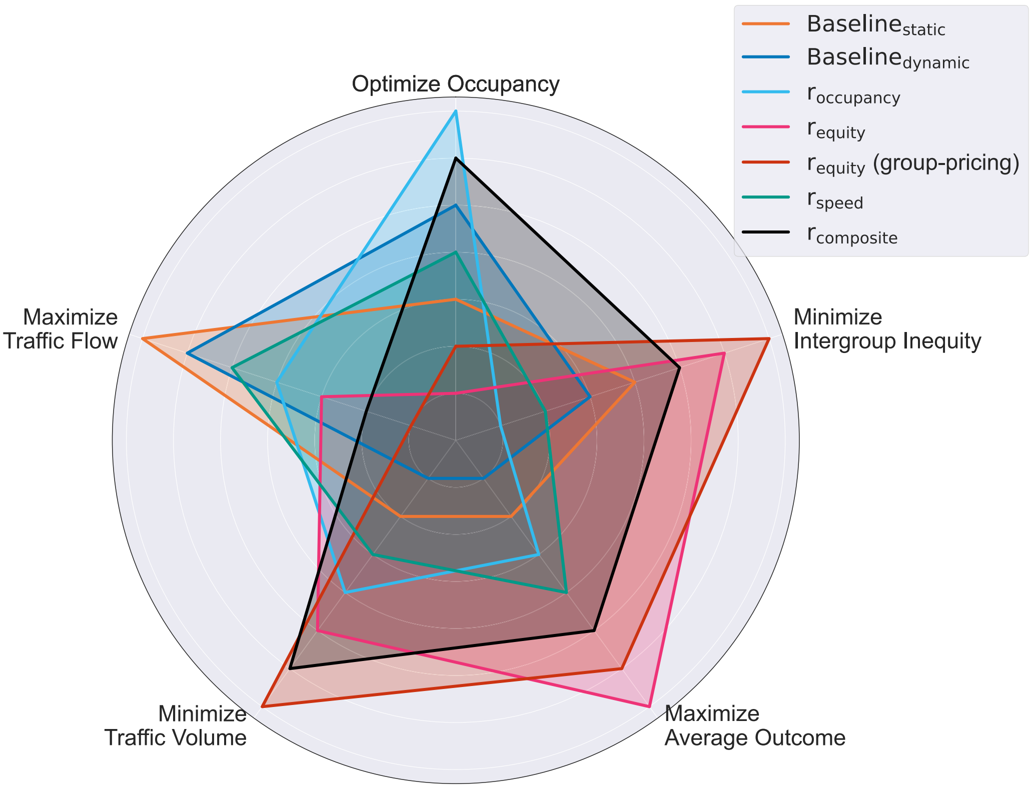

Fig. 7. Rank of the pricing strategies regarding the respective dimension.

Optimizing $r_{occupancy}$, the ML-based pricing scheme posts the best performance regarding managing the parking spaces in the simulation out of all strategies examined, clearly besting the deployed baselines. Nonetheless, the policy learned also results in the highest inequity. Its generally aggressive pricing lowers the outcomes of lower-income cars and renders the one CPZ (teal) it prices affordably so popular that severe congestion ensues in the afternoon, with the average speeds of both through and cruising traffic remaining at close to 5 km/h as all cars attempt to secure the cheap parking spaces. Consequently, even though it succeeded in optimizing its target dimension, the pricing strategy governed by $r_{occupancy}$ does not appear well-suited for everyday use as it illustrates the downsides of exclusively focusing on efficient parking supply management.

The learners attempting to converge towards a fair pricing policy for parking, r_{equity} and r_{equity} with group-specific pricing excel in doing so as they achieve the second-lowest and lowest inequity across all pricing schemes, respectively. While they otherwise perform quite similarly, with unimpressive scores in the occupancy and traffic flow dimensions, r_{equity} also maximizes the average outcome of the agents. This is due to the specific pricing policies chosen by the respective RL agents. The fees set by the system maximizing r_{equity} without group-specific fees essentially correspond to free curbside parking as the RL agent reduces all fees to €0. Since this negates the negative impact of prices on the agents' utilities and outcomes, they are maximized. Crucially, this approach is only able to account for differences in the average utilities of the income groups directly caused by the fees charged. However, there are also indirect effects at play, as, for instance, followers of the parking strategy "car park" are, on average, more affluent and obtain more utility from parking at car parks. Thus, equipped with the ability to set group-specific fees for the different CPZs in the simulation, the ML system can better optimize for r_{equity}. It predominantly selects higher fees for the two higher-income groups, thus better compensating for both the direct and indirect effects of income on the outcomes achieved by drivers by not only increasing the outcomes of lower-income groups but also reducing those of higher-income groups.

r_{speed} exhibits a relatively broad performance profile by posting the third-best scores regarding its target dimension traffic flow (closely behind the baselines), while showing average to good performances in most other dimensions. It achieves this feat by choosing a comparatively aggressive pricing policy that leads to the second-highest level of inequity observed and a sizeable amount of variance in the outcomes produced throughout a modeled day.

Finally, $r_{composite}$ is aimed at combining $r_{occupancy}$, performing best for its target dimension but worst at minimizing inequity, with r_{equity} with group-specific fees, also showing the best result for its targeted metric and delivering the third-worst result regarding the management of the curbside parking supply in the model. The ML-based pricing scheme that optimizes for this combination indeed delivers the second-best occupancy performance and third-best degree of inequity behind the dedicated learners, thus demonstrating that the ML approach allows for the successful combination of different and potentially conflicting goals. It is noteworthy, however, that it does so by converging on relatively erratic pricing policies for the different CPZs consisting of frequent high-interval price jumps.

7. Conclusion

Municipal decision-makers can pursue a variety of policy goals when designing parking pricing strategies. For instance, these policy goals could aim at improving residents' quality of life by achieving the more specific goals of reducing congestion and cumbersome searching times for parking spaces. However, it is crucial to also consider how such policies may affect different subgroups of the population differently and to mitigate potentially inequitable and unintended social outcomes. To research how different parking policies achieve different policy goals including equity between income groups, in the present paper, we first presented a parking simulation featuring an empirically calibrated behavioral model. We then conducted extensive experiments to investigate the fairness of different pricing strategies for parking and evaluate a ML-based system for determining parking fees, optimizing various reward functions corresponding to different goals pursued by municipal decision-makers. In doing so, our core findings are:

Dynamic pricing for parking may negatively impact fairness. While our model, as all parking simulations, cannot represent reality completely accurately (Shoup, 2021), and our dynamic baseline pricing is much more responsive than currently deployed systems, our experiments suggest that the higher fees necessary to conform to such schemes disproportionately impact the outcomes of low-income drivers.

Achieving fair parking may require distinct fees for different social groups. Due to the differences in the outcome structures between income groups, our experiments demonstrated that only group-specific pricing succeeded in reducing inter-group inequity when curbside parking is not offered for free, a scenario which is associated with adverse outcomes for urban societies in general (Shoup, 2011). While such a policy may seem unrealistic for pricing systems used in practice, municipalities have begun experimenting with potential proxy variables, such as vehicle size (Willsher, 2023), to differentiate parking fees.

Pricing parking optimally for individual dimensions comes with unintended consequences. While the ML-based pricing system consistently outperformed our baseline schemes, it also exemplified the pitfalls of focusing on individual policy goals when devising parking fees. Every learner focusing on a dedicated dimension performed poorly in one or multiple of the others, as r_{equity} fails to manage parking supply efficiently, while both $r_{occupancy}$ and r_{speed} result in very high levels of inequity. Consequently, pricing policies for parking should be devised holistically to avoid such unintended consequences.

ML -based pricing strategies offer potential for combining different policy goals. As evidenced by the success of $r_{composite}$, our experiments show that the ML-based pricing system deployed here shows potential to jointly optimize different goals that may seem conflicting at first. Nonetheless, it is important to stress that this approach is only viable for goals to the degree that they are combinable in theory and most likely requires sacrificing performance in the individual dimensions.

Crucially, there are several limitations of our work to note. With respect to survey design, in order to model parking behavior more accurately, instead of forcing them to choose a parking space, it would have been beneficial to include a reject option for respondents in the DCE so as to incorporate the potential rejection of all given alternatives into our simulation.

Regarding the simulation, while we strived to calibrate it with empirical data, it is worth pointing out that our ABM does not achieve the same level of spatial accuracy when compared with existing models based on GIS data. Crucially, our approach of simulating individual days limits the degree to which drivers can respond to the pricing strategies tested by, e.g., retiming their trips to anticipate and avoid times high traffic volume and parking usage. Moreover, the parking simulation does not account for the level of information available to drivers. Future research could add further extensions that let the available information vary, include illegal parking behavior and expand the simulated time frame. Furthermore, the simulation could be extended with respect to further individual-level attributes to inspect not only inequities with respect to groups but also, for instance, car emissions.

Finally, it remains to acknowledge that the ML-based pricing system established in this study is, of course, quite far away from representing a realistic option for practitioners. Shoup (2021) lists nine requirements for pricing schemes for curbside parking. While for some of them, the policies examined during the experiments represent a good fit, they fail to meet certain requirements, such as transparency or stability of fees, thus negating advantages they may bring when it comes to the fairness requirement. Nonetheless, we believe our findings suggest that such systems provide interesting insights into pricing policies necessary to achieve specific policy goals and pricing parking both fairly and in accordance with the social costs incurred. Given some restrictions, e.g., on the variability of fees chosen, a similar approach may become more viable in the future. This holds especially as municipalities increasingly attempt to harness advancements in information and communication technologies to allow for extensive data collection and processing as well as the use of said data to automate and predict usage of city services to become "smart cities" (Dustdar et al., 2017). Apart from addressing these limitations, it also remains for future research to investigate how theoretically income-agnostic policies, such as banning all cars from the city center, affect the outcomes achieved by different social groups when venturing into cities.

From a broader perspective, our experiments demonstrate the viability of using simulations to evaluate the impact and fairness of public policies in urban contexts. This holds in particular since public policies in general and those relying on new methodologies, such as ML, in particular, may bring wide-ranging consequences that can only be partially understood beforehand. Moreover, the conceptualization of fairness developed for this study allows for devising policies that promote effectiveness and efficiency while adhering to fairness constraints. Even though the scenario presented is significantly simplified, with equity modeled solely with regard to income, the approach adopted here can, in principle, be applied to more complex settings.

CRediT authorship contribution statement

Jakob Kappenberger: Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Formal analysis, Conceptualization. **Heiner Stuckenschmidt:** Writing – review & editing, Writing – original draft, Conceptualization. **Frederic Gerdon:** Writing – review & editing, Investigation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research was conducted as part of the grant "Consequences of Artificial Intelligence for Urban Societies (CAIUS), Germany", funded by Volkswagen Foundation. We want to extend our thanks to our colleague Daria Szafran for her work in conceptualizing and creating the survey that forms the basis of our behavioral model. Furthermore, we would like to thank our colleagues Lea Cohausz, Ruben Bach, Florian Rupp, and Christoph Kern as well as the reviewers for their helpful feedback on an earlier version of this article.

Appendix A. Methodology

A.1. Survey experiment design

Table A.3 shows the attributes and their respective levels chosen for the survey experiment. The attributes were selected in accordance with previous studies (see Section 3.1) to cover a realistic range of individual values and limit the number of resulting choice sets by, e.g., focusing on two major types of parking spaces in city centers: curbside parking and car parks. When determining the choice sets (i.e., the choice tasks combining the different levels of the attributes), one attempts to strike a balance between the number of choice sets (more choice sets will lead to a smaller number of respondents per choice set) and reducing the variance of the coefficients estimated based on the design (usually measured by the *d-error*) (Hensher et al., 2015). Since the full factorial of these attributes (i.e., the choice sets containing all potential alternatives) would amount to 216 choice sets, we opted for a more efficient approach by deploying the *idefix* package for R (R Core Team, 2021), which allows for generating designs that minimize the d-error and incorporating a prior distribution of the coefficients of the different attributes (Traets et al., 2020). We supplied the results reported by Axhausen and Polak (1991) as our prior distribution, as their survey was conducted in a similar setting, albeit over 30 years earlier. This procedure resulted in 64 choice sets utilized for the survey.

For these 64 sets, we automatically generated visualizations, such as the one shown in Fig. A.8, to ease the choice tasks for respondents. Additionally, every choice task was accompanied by a textual prompt that contained different values for the aforementioned variables *trip purpose* (either "Work/Education", "Doctor's appointment", Meeting an Acquaintance", or "Shopping") and *time of day* (either "morning", "midday", or "afternoon").[15] As an example, the resulting text for the *trip purpose* "Doctor's appointment" at the *time of day* "morning" would read:

"Imagine heading to the nearest city by car on a business day **in the morning** for **a doctor's appointment**. Your destination is marked with a "X" on the graphic below. You can choose between two parking spaces (described in the red and blue boxes). Which one do you prefer?" (translated from German)

Every respondent was asked to complete four random choice tasks so as not to overburden the respondents, given that the DCE only formed a portion of the overall survey. Additionally, respondents were asked which of the parking strategies (see Section 3.1) best described their general approach to parking search.

[15] These options represent a simplified version of those synthesized in Anon (2018).

Table A.3
Attributes and attribute levels for the DCE.

Attribute	Levels
Access	5 min, 10 min, or 15 min
Search Time	2 min, 5 min, or 10 min
Egress	2 min, 5 min, or 10 min
Fee	€0, €1, €2, or €4
Space Type	Curbside or car park

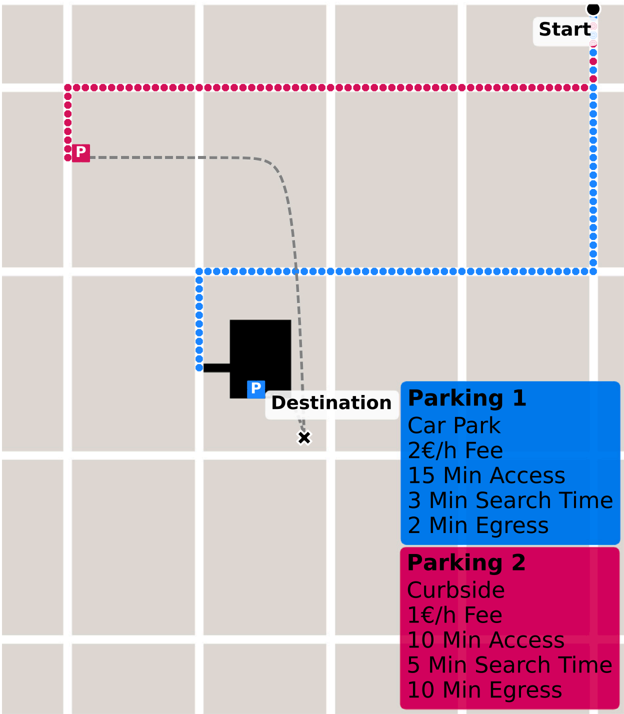

Fig. A.8. Map visualizing two parking options for survey respondents (translated from German).

Table A.4
Distribution of Age and Gender in Survey Sample.

Age Group	Female	Male	Diverse
18–36	248	271	5
37–57	318	307	1
58–80	398	326	1

A.2. Survey sample and logit models

See Tables A.4–A.7.

Appendix B. Simulation

See Figs. B.9 and B.10.

B.1. Model calibration

We created two different versions of the ABM for parking, one for training the ML-based pricing system (see Section 6.1) and one for the overall evaluation of all schemes. For the former, shorter computation times are crucial to achieving the number of episodes required for the type of ML deployed (see Section 6.1). This holds in particular for our case since the complexity of the underlying parking simulation, with its large number of agents and subroutines, results in relatively long runtimes, even without accounting for computing overhead of ML. Using a larger scale evaluation model (the one shown in Fig. 3), both in spatial terms and regarding

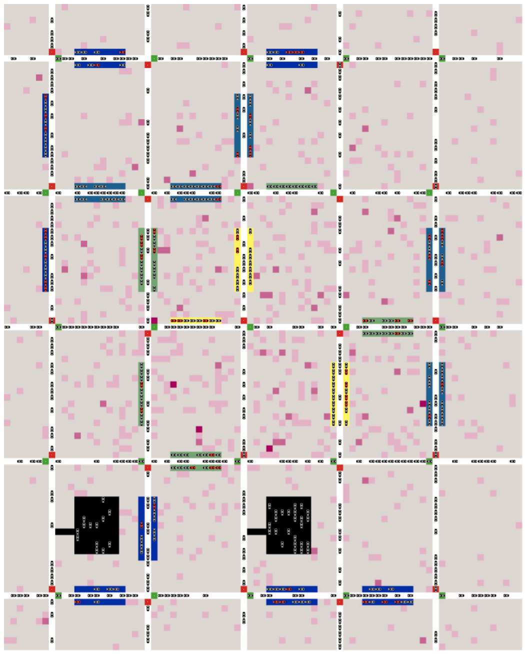

Fig. B.9. Distribution of destinations (magenta points) in simulation. Darker hues indicate more destinations at an individual point.

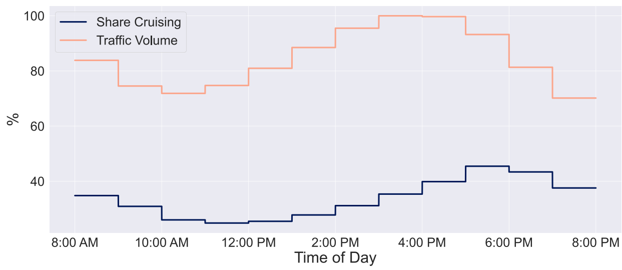

Fig. B.10. Distributions of relative traffic volume and share of traffic cruising for parking during static baseline run.

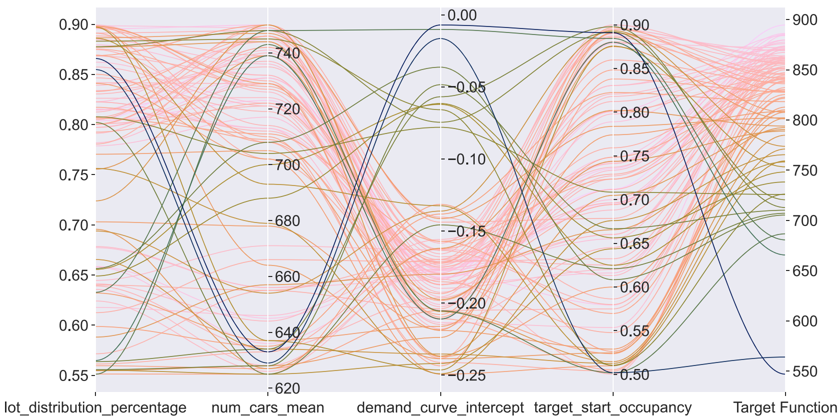

Fig. B.11. Result of Bayesian hyperparameter tuning over 100 episodes for the evaluation model. The target function optimizes for the desired volume of traffic, an average speed of 15 km/h, and an average share of cruising cars of 30% (consistent with Hampshire and Shoup (2019)) (higher is better).

Table A.5
Sample Descriptive Statistics.

Variable	Value	Percentage
Gender	Female	48.8%
	Male	51.2%
Age	18–36	28.2%
	37–57	33.9%
	58–80	37.5%
Household Income	<€520	0.8%
	€520–€750	0.4%
	€750–€1,500	0.6%
	€1,500–€2,500	16.3%
	€2,500–€3,500	24.0%
	€3,500–€5,000	29.0%
	>€5,000	23.9%
Strategy	Close to Goal	34.7%
	Car Park	28.8%
	En Route	19.3%
	Other	17.2%

Table A.6
Mixed Logit Model of Parking Choice (only Main Effects). σ refers to the standard deviation of the coefficient mean if random coefficients are estimated.

	Coefficient		σ	
Access (min)	−0.10	(0.01)***	0.03	(0.05)
Search Time (min)	−0.12	(0.01)***	0.16	(0.04)***
Egress (min)	−0.26	(0.02)***	0.22	(0.03)***
Space Type Car Park	0.53	(0.07)***	0.99	(0.18)***
Fee (€)	−1.44	(0.10)***	0.88	(0.08)***
Age	−0.00	(0.00)	0.01	(0.00)**
Gender Female	0.18	(0.10)*	0.43	(0.28)
R^2_{McFadden}	0.33			
Log Likelihood	−2920.56			
Num. obs.	6301			

$p < 0.01$; $p < 0.05$;* $p < 0.1$

Table A.7
Complete Mixed Logit Model of Parking Choice. σ refers to the standard deviation of the coefficient mean if random coefficients are estimated.

	Coefficient		σ	
Access (min)	−0.04	(0.02)*	−0.04	(0.04)
Search Time (min)	−0.05	(0.04)	0.14	(0.04)***
Egress (min)	−0.24	(0.04)***	0.20	(0.03)***
Space Type Car Park	−0.05	(0.18)	0.80	(0.19)***
Fee (€)	−1.23	(0.26)***	0.84	(0.08)***
Age	−0.00	(0.00)	0.01	(0.00)**
Gender Female	0.16	(0.10)*	0.34	(0.33)
Strategy En Route * Access	0.00	(0.03)	–	
Strategy En Route * Search Time	−0.03	(0.04)	–	
Strategy En Route * Egress	0.10	(0.03)***	–	
Strategy En Route * Space Type Car Park	−0.18	(0.18)	–	
Strategy En Route * Fee	−0.01	(0.09)	–	
Strategy Car Park * Access	0.01	(0.02)	–	
Strategy Car Park * Search Time	−0.05	(0.03)	–	
Strategy Car Park * Egress	0.04	(0.03)	–	
Strategy Car Park * Space Type Car Park	0.87	(0.16)***	–	
Strategy Car Park * Fee	0.52	(0.08)***	–	
Strategy Other * Access	−0.04	(0.03)	–	
Strategy Other * Search Time	−0.09	(0.04)**	–	
Strategy Other * Egress	0.12	(0.04)***	–	
Strategy Other * Space Type Car Park	−0.07	(0.18)	–	
Strategy Other * Fee	−0.11	(0.10)	–	
Time Midday * Access	−0.02	(0.02)	–	

(continued on next page)

Table A.7 (*continued*).

Time Midday * Search Time	0.03	(0.03)	–
Time Midday * Egress	0.01	(0.03)	–
Time Midday * Space Type Car Park	0.19	(0.15)	–
Time Midday * Fee	0.10	(0.07)	–
Time Afternoon * Access	−0.03	(0.02)	–
Time Afternoon * Search Time	−0.02	(0.03)	–
Time Afternoon * Egress	0.02	(0.03)	–
Time Afternoon * Space Type Car Park	0.25	$(0.15)^{*}$	–
Purpose Doctor * Access	−0.03	(0.02)	–
Purpose Doctor * Search Time	−0.07	$(0.04)^{*}$	–
Purpose Doctor * Egress	−0.09	$(0.03)^{***}$	–
Purpose Doctor * Space Type Car Park	0.17	(0.17)	–
Purpose Doctor * Fee	0.68	$(0.09)^{***}$	–
Purpose Acquaintance * Access	−0.08	$(0.03)^{***}$	–
Purpose Acquaintance * Search Time	−0.02	(0.04)	–
Purpose Acquaintance * Egress	−0.09	$(0.03)^{**}$	–
Purpose Acquaintance * Space Type Car Park	0.35	$(0.18)^{*}$	–
Purpose Acquaintance * Fee	0.24	$(0.09)^{***}$	–
Purpose Shopping * Access	−0.04	(0.02)	–
Purpose Shopping * Search Time	−0.05	(0.04)	–
Purpose Shopping * Egress	−0.15	$(0.03)^{***}$	–
Purpose Shopping * Space Type Car Park	0.31	$(0.17)^{*}$	–
Purpose Shopping * Fee	0.52	$(0.09)^{***}$	–
Household Income 2 * Fee	−1.01	$(0.48)^{**}$	–
Household Income 3 * Fee	−0.74	$(0.27)^{***}$	–
Household Income 4 * Fee	−0.86	$(0.25)^{***}$	–
Household Income 5 * Fee	−0.75	$(0.25)^{***}$	–
Household Income 6 * Fee	−0.74	$(0.25)^{***}$	–
Household Income 7 * Fee	−0.67	$(0.25)^{***}$	–
R^2_{McFadden}	0.37		
Log Likelihood	−2764.19		
Num. obs.	6301		

$^{**}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

the number of agents simulated, we expect a better representation of our empirical calibration and, thus, more conservative and robust results.

To obtain a representative model configuration, we conducted hyperparameter tuning[16] to calibrate the simulation based on government data provided by our partnering municipality: The number of 47,000 cars entering and leaving the modeled area was translated to the 12-hour period simulated and the number of blocks in the ABM. The model was then run in its static baseline configuration (described in Section 6.1) with different values for its initial hyperparameters to achieve both a representative overall traffic volume and flow (measured as approx. 15 km/h in the modeled area of Mannheim) as well as a share of cruising cars consistent with studies related to our use case (Hampshire and Shoup, 2019).[17] Fig. B.11 serves to visualize this process. For the smaller training configuration, we then searched for the set of parameters conducive to creating a simulation that was as similar as possible to our evaluation model (Fig. B.12 shows the different configurations tested). The resulting parameter configurations can be inspected in Appendix C.

Appendix C. Parameters for reproduction

See Tables C.8–C.10.

Appendix D. Detailed results

The following plots describe the results of our experiments in more detail. For every pricing strategy, we selected the run with the median accumulated score over the evaluation period (see Figs. D.13–D.19).

[16] All hyperparameter tuning runs and experiments featured in this paper were conducted using *Weights and Biases* (Biewald, 2020).

[17] It is important to note that the frequently cited average share of cruising cars of 30% has to be used cautiously, as different studies report wildly different values (Shoup, 2021). Due to a lack of further empirical evidence, we still had to rely on this convention as an average over multiple simulations.

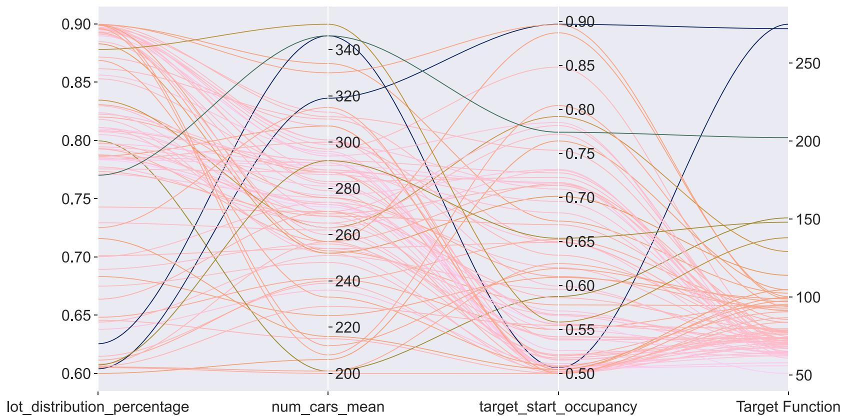

Fig. B.12. Result of random hyperparameter tuning for 100 episodes for the training model. The target function minimizes the Euclidean distance to the evaluation configuration in terms of average occupancy, traffic volume, and flow (lower is better).

Table C.8
Model parameters of our NetLogo ABM.

Variable	Value
ticks	21,600
num-cars-mean	359 (training), 800 (evaluation)
max-x-cor	28 (training), 40 (evaluation)
max-y-cor	30 (training), 50 (evaluation)
num-garages	1 (training), 2 (evaluation)
count-curb-spaces	150 (training), 330 (evaluation)
count-garage-spaces	63 (training), 126 (evaluation)
lot-distribution-percentage	0.60 (training), 0.62 (evaluation)
target-start-occupancy	0.51 (training), 60 (evaluation)
demand-curve-intercept	0.25
initial fee of all CPZs	€2.0
pop-mean-income	€3,612
pop-median-income	€2,956
temporal-resolution	1,800

Table C.9
Variables contained in state representation S_t for RL.

Variable	Description
ticks	Ticks passed so far in the simulation
n_cars	Share of originally spawned vehicles currently in simulation
mean_speed	Normalized average speed of all non-parking cars
CPZ_occupancy	Current utilized capacities of all four CPZs
garages_occupancy	Occupancy of garage(s) (average if multiple)
global_inequity	Individual Inequity
intergroup_inequity	Inter-group Inequity

Table C.10
Final PPO hyperparameters across all reward functions.

Function	Batch size	N steps	Gamma	Learning rate
$\mathrm{r}_{\mathrm{occupancy}}$	3072	48	0.999	0.0016
$\mathrm{r}_{\mathrm{equity}}$	7680	24	0.9	0.0003
$\mathrm{r}_{\mathrm{equity}(group)}$	3072	48	0.99	0.0005
$\mathrm{r}_{\mathrm{speed}}$	12 288	48	0.995	0.0012
$\mathrm{r}_{\mathrm{composite}(group)}$	3072	48	0.99	0.0005

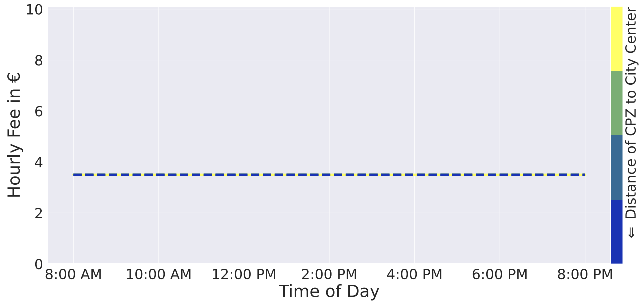

(a) CPZ fees.

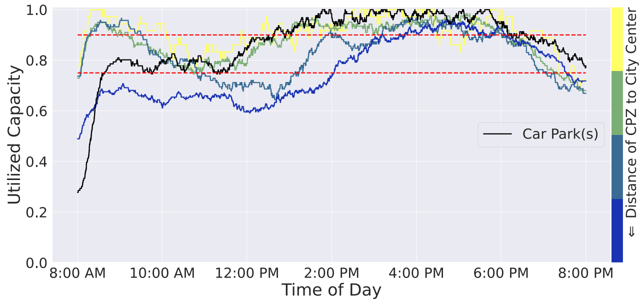

(b) Occupancy levels of CPZs with desired corridor in red.

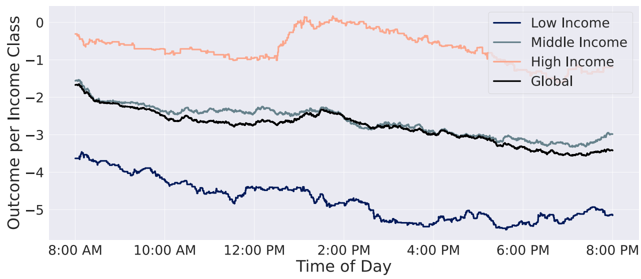

(c) Outcomes of different income groups.

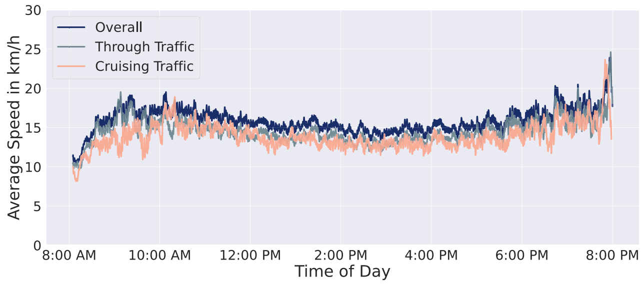

(d) Average speed of unparked cars.

Fig. D.13. Results of the median static baseline run.

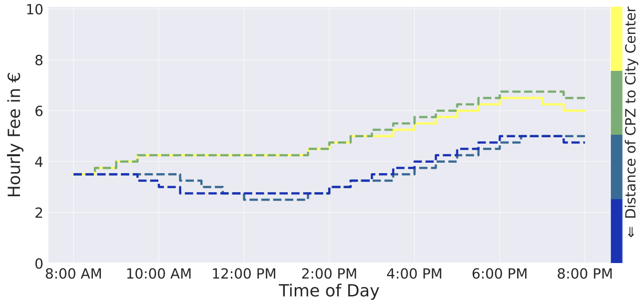

(a) CPZ fees.

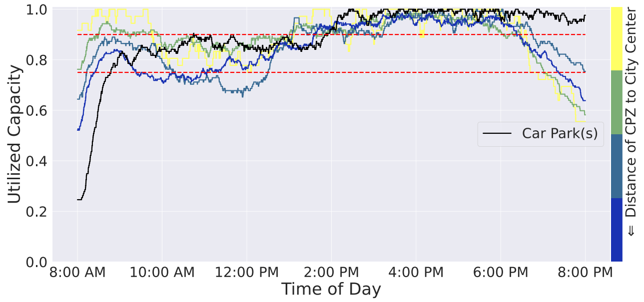

(b) Occupancy levels of CPZs with desired corridor in red.

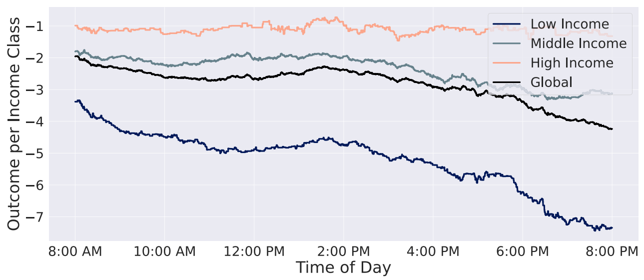

(c) Outcomes of different income groups.

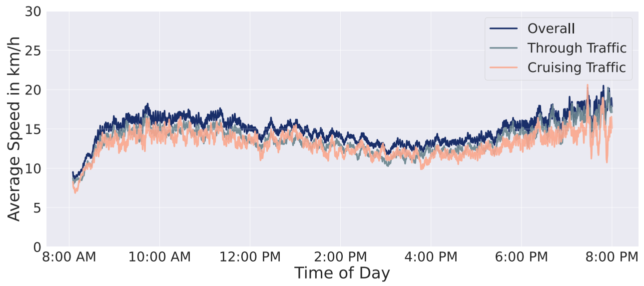

(d) Average speed of unparked cars.

Fig. D.14. Results of the median dynamic baseline run.

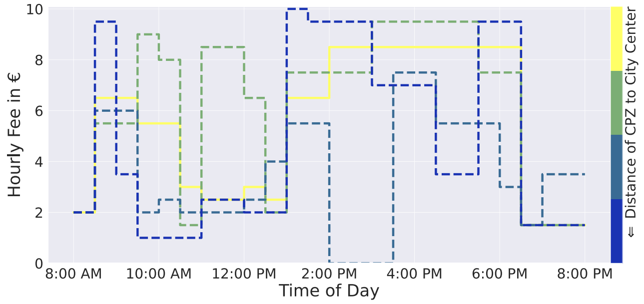

(a) CPZ fees.

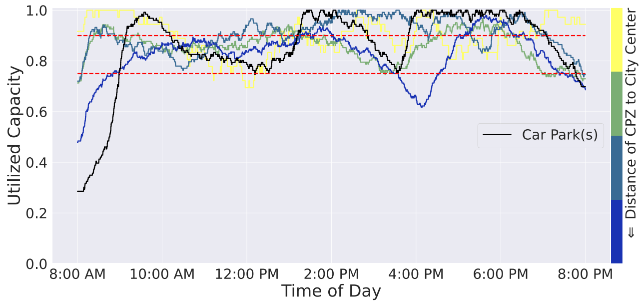

(b) Occupancy levels of CPZs with desired corridor in red.

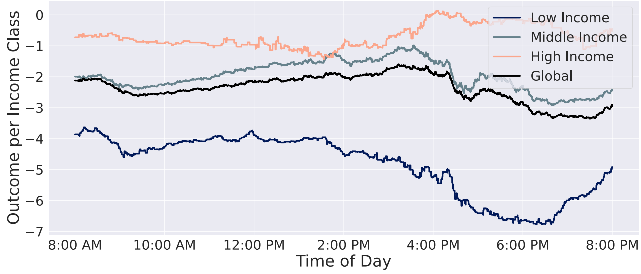

(c) Outcomes of different income groups.

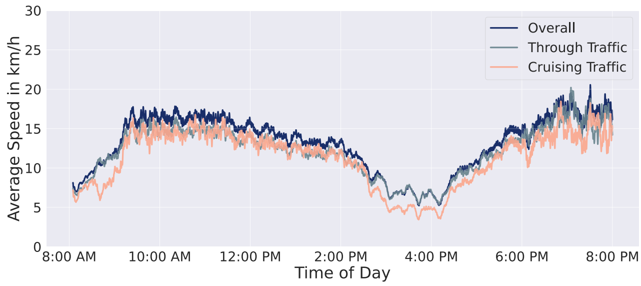

(d) Average speed of unparked cars.

Fig. D.15. Results of the median $r_{occupancy}$ run.

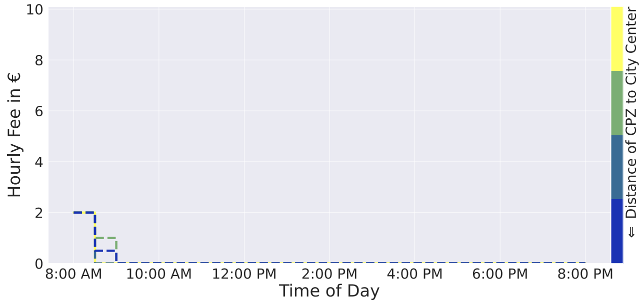

(a) CPZ fees.

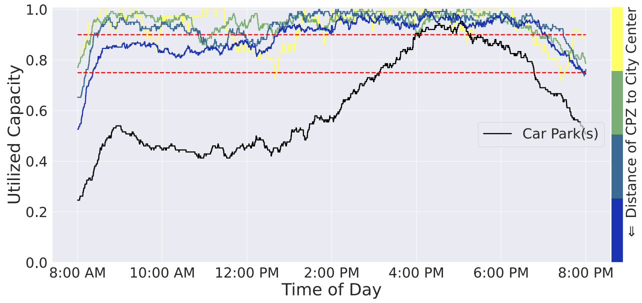

(b) Occupancy levels of CPZs with desired corridor in red.

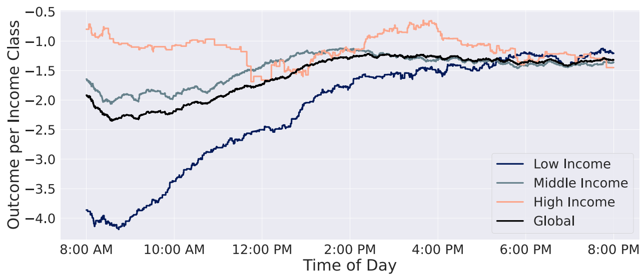

(c) Outcomes of different income groups.

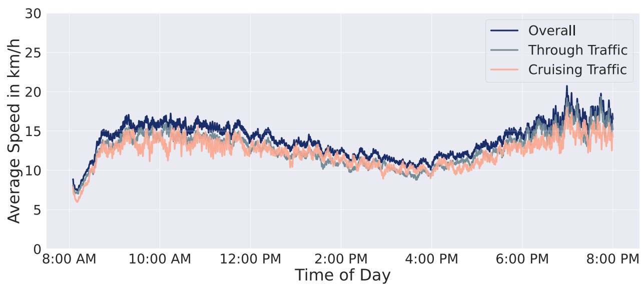

(d) Average speed of unparked cars.

Fig. D.16. Results of the median r_{equity} run.

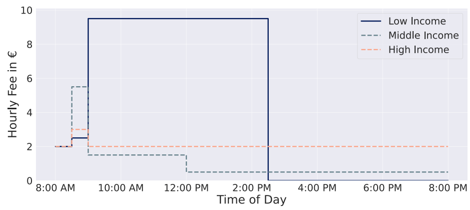

(a) Group-specific fees for the yellow CPZ.

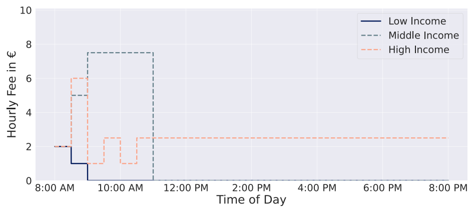

(b) Group-specific fees for the green CPZ.

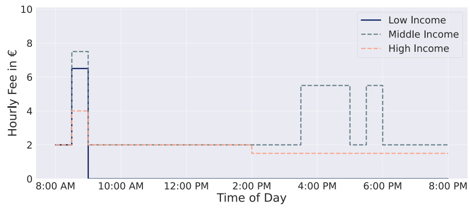

(c) Group-specific fees for the teal CPZ.

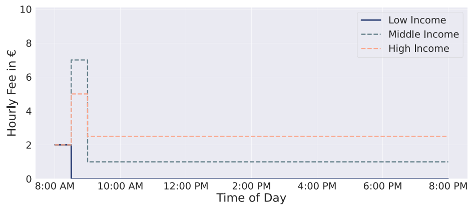

(d) Group-specific fees for the blue CPZ.

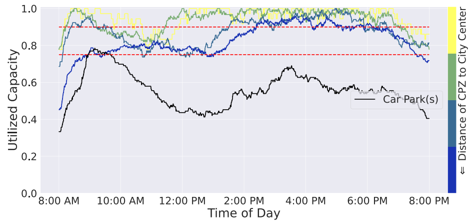

(e) Occupancy levels of CPZs with desired corridor in red.

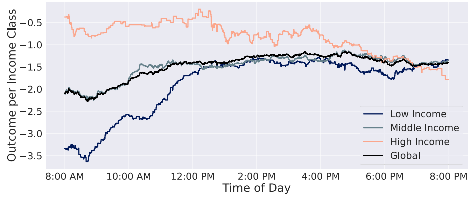

(f) Outcomes of different income groups.

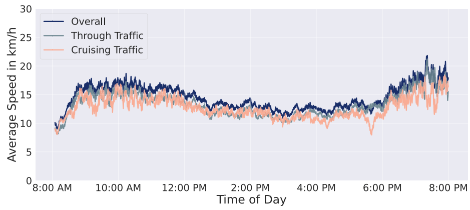

(g) Average speed of unparked cars.

Fig. D.17. Results of the median r_{equity} run with group-specific pricing.

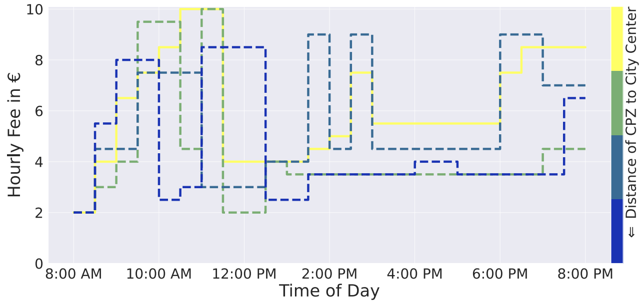

(a) CPZ fees.

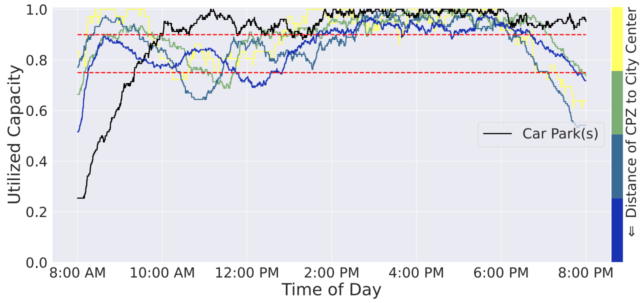

(b) Occupancy levels of CPZs with desired corridor in red.

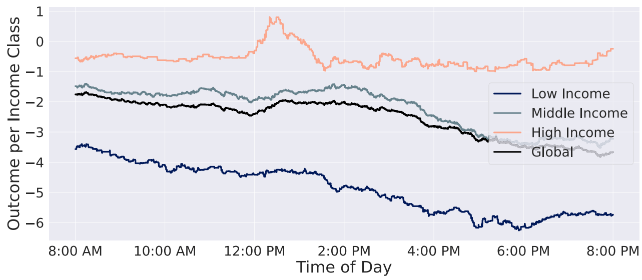

(c) Outcomes of different income groups.

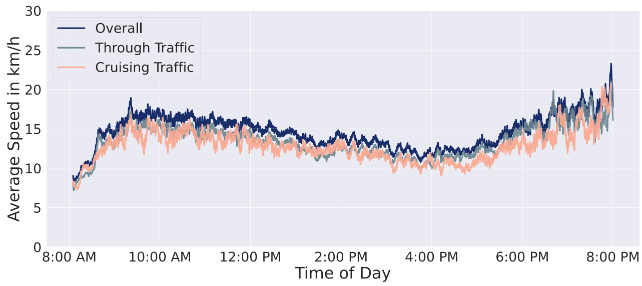

(d) Average speed of unparked cars.

Fig. D.18. Results of the median r_{speed} run.

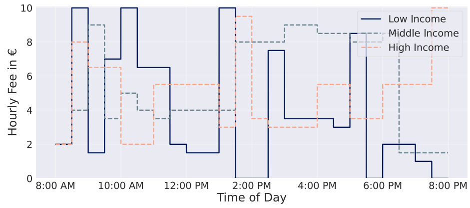

(a) Group-specific fees for the yellow CPZ.

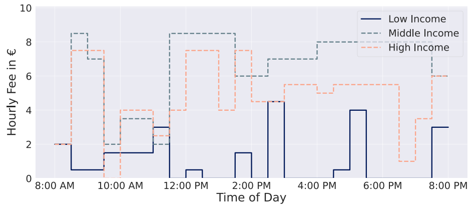

(b) Group-specific fees for the green CPZ.

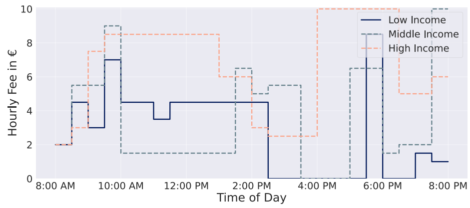

(c) Group-specific fees for the teal CPZ.

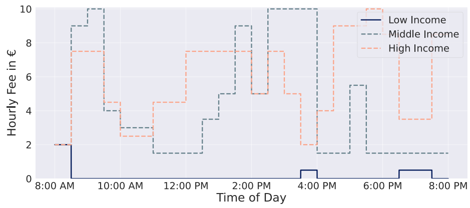

(d) Group-specific fees for the blue CPZ.

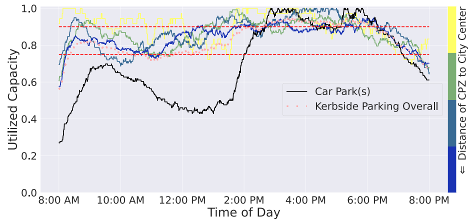

(e) Occupancy levels of CPZs with desired corridor in red.

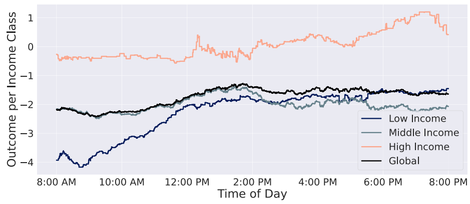

(f) Outcomes of different income groups.

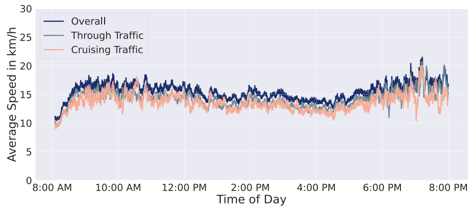

(g) Average speed of unparked cars.

Fig. D.19. Results of the median $r_{composite}$ run with group-specific pricing.

References

Anon, 2018. Mobilität in Deutschland. Technical Report, infas, DLR, IVT, and infas 360, URL: https://www.mobilitaet-in-deutschland.de/archive/pdf/MiD2017_Ergebnisbericht.pdf.

Antolin, G., Ibeas, A., Alonso, B., dell'Olio, L., 2018. Modelling parking behaviour considering users heterogeneities. Transp. Policy 67, 23–30. http://dx.doi.org/10.1016/j.tranpol.2018.01.014, URL: https://linkinghub.elsevier.com/retrieve/pii/S0967070X16304334.

Arnott, R., de Palma, A., Lindsey, R., 1991. A temporal and spatial equilibrium analysis of commuter parking. J. Public Econ. 45 (3), 301–335. http://dx.doi.org/10.1016/0047-2727(91)90030-6, URL: https://linkinghub.elsevier.com/retrieve/pii/0047272791900306.

Axhausen, K.W., Polak, J.W., 1991. Choice of parking: Stated preference approach. Transportation 18 (1), 59–81. http://dx.doi.org/10.1007/BF00150559, URL: http://link.springer.com/10.1007/BF00150559.

Barocas, S., Hardt, M., Narayanan, A., 2019. Fairness and Machine Learning. fairmlbook.org, http://www.fairmlbook.org.

Benenson, I., Martens, K., Birfir, S., 2008. PARKAGENT: An agent-based model of parking in the city. Comput. Environ. Urban Syst. 32 (6), 431–439. http://dx.doi.org/10.1016/j.compenvurbsys.2008.09.011, URL: https://linkinghub.elsevier.com/retrieve/pii/S0198971508000689.

Biewald, L., 2020. Experiment tracking with weights and biases. URL: https://www.wandb.com/. Software available from wandb.com.

Brooke, S., Ison, S., Quddus, M., 2014. On-street parking search: Review and future research direction. Transp. Res. Record: J. Transp. Res. Board 2469 (1), 65–75. http://dx.doi.org/10.3141/2469-08, URL: http://journals.sagepub.com/doi/10.3141/2469-08.

Cascetta, E., 2009. Random utility theory. In: Transportation Systems Analysis, vol. 29, Springer US, Boston, MA, pp. 89–167. http://dx.doi.org/10.1007/978-0-387-75857-2_3, URL: http://link.springer.com/10.1007/978-0-387-75857-2_3. Series Title: Springer Optimization and Its Applications.

Chaniotakis, E., Pel, A.J., 2015. Drivers' parking location choice under uncertain parking availability and search times: A stated preference experiment. Transp. Res. A: Policy Pr. 82, 228–239. http://dx.doi.org/10.1016/j.tra.2015.10.004, URL: https://linkinghub.elsevier.com/retrieve/pii/S0965856415002542.

Croissant, Y., 2020. Estimation of random utility models in R: The mlogit package. J. Stat. Softw. 95 (11), 1–41. http://dx.doi.org/10.18637/jss.v095.i11.

Deutsch, M., 1975. Equity, equality, and need: What determines which value will be used as the basis of distributive justice? J. Soc. Issues 31 (3), 137–149. http://dx.doi.org/10.1111/j.1540-4560.1975.tb01000.x, URL: https://spssi.onlinelibrary.wiley.com/doi/10.1111/j.1540-4560.1975.tb01000.x.

Dieussaert, K., Aerts, K., Steenberghen, T., Maerivoet, S., Spitaels, K., 2009. SUSTAPARK: An agent-based model for simulating parking search. In: 12th AGILE International Conference on Geographic Information Science Hannover. Hannover, pp. 1–11.

Domencich, T.A., McFadden, D., 1975. Urban travel demand: A behavioral analysis. In: Contributions to economic analysis, (93), North-Holland Pub. Co. ; American Elsevier, Amsterdam, New York.

Dustdar, S., Nastić, S., Šćekić, O., 2017. Smart Cities. Springer International Publishing, Cham, http://dx.doi.org/10.1007/978-3-319-60030-7, URL: http://link.springer.com/10.1007/978-3-319-60030-7.

Ellis, R.H., Bennett, J.C., Rassam, P.R., 1974. Parking systems analysis: an overview. ACM SIGSIM Simul. Dig. 6 (1), 42–50. http://dx.doi.org/10.1145/1102934.1102941, URL: https://dl.acm.org/doi/10.1145/1102934.1102941.

Friedel, J.E., Foreman, A.M., Wirth, O., 2022. An introduction to "discrete choice experiments" for behavior analysts. Behav. Process. 198, 104628. http://dx.doi.org/10.1016/j.beproc.2022.104628, URL: https://linkinghub.elsevier.com/retrieve/pii/S0376635722000456.

Friesen, M., Mingardo, G., 2020. Is parking in Europe ready for dynamic pricing? A reality check for the private sector. Sustainability 12 (7), 2732. http://dx.doi.org/10.3390/su12072732, URL: https://www.mdpi.com/2071-1050/12/7/2732.

Grgic-Hlaca, N., Zafar, M.B., Gummadi, K.P., Weller, A., 2016. The case for process fairness in learning: Feature selection for fair decision making. In: NIPS Symposium on Machine Learning and the Law, Vol. 1. p. 2.

Hampshire, R., Shoup, D., 2019. How much traffic is cruising for parking? Transf. Mag. 4, URL: https://transfersmagazine.org/magazine-article/issue-4/how-much-traffic-is-cruising-for-parking/.

Hensher, D.A., Rose, J.M., Greene, W.H., 2015. Applied Choice Analysis, second ed. Cambridge University Press, http://dx.doi.org/10.1017/CBO9781316136232, URL: https://www.cambridge.org/core/product/identifier/9781316136232/type/book.

Hess, S., Daly, A., Batley, R., 2018. Revisiting consistency with random utility maximisation: theory and implications for practical work. Theory and Decision 84 (2), 181–204. http://dx.doi.org/10.1007/s11238-017-9651-7, URL: http://link.springer.com/10.1007/s11238-017-9651-7.

Hilvert, O., Toledo, T., Bekhor, S., 2012. Framework and model for parking decisions. Transp. Res. Record: J. Transp. Res. Board 2319 (1), 30–38. http://dx.doi.org/10.3141/2319-04, URL: http://journals.sagepub.com/doi/10.3141/2319-04.

Horni, A., Montini, L., Waraich, R.A., Axhausen, K.W., 2013. An agent-based cellular automaton cruising-for-parking simulation. Transp. Lett. 5 (4), 167–175. http://dx.doi.org/10.1179/1942787513Y.0000000004, URL: http://www.tandfonline.com/doi/full/10.1179/1942787513Y.0000000004.

Inci, E., 2015. A review of the economics of parking. Econ. Transp. 4 (1–2), 50–63. http://dx.doi.org/10.1016/j.ecotra.2014.11.001, URL: https://linkinghub.elsevier.com/retrieve/pii/S2212012214000562.

Int Panis, L., Broekx, S., Liu, R., 2006. Modelling instantaneous traffic emission and the influence of traffic speed limits. Sci. Total Environ. 371 (1–3), 270–285. http://dx.doi.org/10.1016/j.scitotenv.2006.08.017, URL: https://linkinghub.elsevier.com/retrieve/pii/S004896970600636X.

Jakob, M., Menendez, M., 2021. Parking pricing vs. congestion pricing: a macroscopic analysis of their impact on traffic. Transp. A: Transp. Sci. 17 (4), 462–491. http://dx.doi.org/10.1080/23249935.2020.1797924, URL: https://www.tandfonline.com/doi/full/10.1080/23249935.2020.1797924.

Jakob, M., Menendez, M., Cao, J., 2020. A dynamic macroscopic parking pricing and decision model. Transp. B: Transp. Dyn. 8 (1), 307–331. http://dx.doi.org/10.1080/21680566.2018.1488226, URL: https://www.tandfonline.com/doi/full/10.1080/21680566.2018.1488226.

Jaxa-Rozen, M., Kwakkel, J.H., 2018. PyNetLogo: Linking NetLogo with Python. J. Artif. Soc. Soc. Simul. 21 (2), 4. http://dx.doi.org/10.18564/jasss.3668, URL: http://jasss.soc.surrey.ac.uk/21/2/4.html.

Kappenberger, J., Theil, K., Stuckenschmidt, H., 2022. Evaluating the impact of AI-based priced parking with social simulation. In: Hopfgartner, F., Jaidka, K., Mayr, P., Jose, J., Breitsohl, J. (Eds.), In: Social Informatics, vol. 13618, Springer International Publishing, Cham, pp. 54–75. http://dx.doi.org/10.1007/978-3-031-19097-1_4, URL: https://link.springer.com/10.1007/978-3-031-19097-1_4. Series Title: Lecture Notes in Computer Science.

Levy, N., Martens, K., Benenson, I., 2013. Exploring cruising using agent-based and analytical models of parking. Transp. A: Transp. Sci. 9 (9), http://dx.doi.org/10.1080/18128602.2012.664575.

Litman, T., 2019. Parking Management Best Practices. Routledge, New York, OCLC: 1110487321.

Macea, L.F., Serrano, I., Carcache-Guas, C., 2023. A reservation-based parking behavioral model for parking demand management in urban areas. Socio- Econ. Plan. Sci. 86, 101477. http://dx.doi.org/10.1016/j.seps.2022.101477, URL: https://linkinghub.elsevier.com/retrieve/pii/S0038012122002786.

Maxner, T., Ranjbari, A., Dowling, C.P., Günes, S., 2023. Simulation-based analysis of different curb space type allocations on curb performance. Transp. B: Transp. Dyn. 11 (1), 1384–1405. http://dx.doi.org/10.1080/21680566.2023.2212324, URL: https://www.tandfonline.com/doi/full/10.1080/21680566.2023.2212324.

McFadden, D., 1979. Quantitative methods for analysing travel behaviour of individuals - Some recent developments. In: Hensher, D.A., Stopher, P.R. (Eds.), Behavioural Travel Modelling, first ed. Routledge, http://dx.doi.org/10.4324/9781003156055, URL: https://www.taylorfrancis.com/books/9781000362978.

Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., Galstyan, A., 2022. A survey on bias and fairness in machine learning. ACM Comput. Surv. 54 (6), 1–35. http://dx.doi.org/10.1145/3457607, URL: https://dl.acm.org/doi/10.1145/3457607.

Menéndez, M., Pardo, J., Pardo, L., Pardo, M., 1997. The Jensen-Shannon divergence. J. Franklin Inst. 334 (2), 307–318. http://dx.doi.org/10.1016/S0016-0032(96)00063-4, URL: https://linkinghub.elsevier.com/retrieve/pii/S0016003296000634.

Mingardo, G., van Wee, B., Rye, T., 2015. Urban parking policy in Europe: A conceptualization of past and possible future trends. Transp. Res. A: Policy Pr. 74, 268–281. http://dx.doi.org/10.1016/j.tra.2015.02.005, URL: https://linkinghub.elsevier.com/retrieve/pii/S0965856415000221.

Montini, L., Horni, A., Rieser-Schüssler, N., Axhausen, K.W., 2012. Searching for parking in GPS data. p. 26. http://dx.doi.org/10.3929/ETHZ-A-007319185, URL: http://hdl.handle.net/20.500.11850/49659. Artwork Size: 26 p. Medium: application/pdf Publisher: [object Object].

OECD, 2019. Under Pressure: The Squeezed Middle Class. OECD, http://dx.doi.org/10.1787/689afed1-en, URL: https://www.oecd-ilibrary.org/social-issues-migration-health/under-pressure-the-squeezed-middle-class_689afed1-en.

OECD, 2023. Passenger cars per one thousand inhabitants. URL: https://stats.oecd.org/Index.aspx?DataSetCode=ITF_INDICATORS#.

Pierce, G., Shoup, D., 2013. Getting the prices right: An evaluation of pricing parking by demand in San Francisco. J. Am. Plan. Assoc. 79 (1), 67–81. http://dx.doi.org/10.1080/01944363.2013.787307, URL: http://www.tandfonline.com/doi/abs/10.1080/01944363.2013.787307.

Ping, H., Stoyanovich, J., Howe, B., 2017. DataSynthesizer: Privacy-preserving synthetic datasets. In: Proceedings of the 29th International Conference on Scientific and Statistical Database Management. ACM, Chicago IL USA, pp. 1–5. http://dx.doi.org/10.1145/3085504.3091117, URL: https://dl.acm.org/doi/10.1145/3085504.3091117.

Polak, J., Axhausen, K., 1990. Parking search behaviour: A review of current research and future prospects. p. 30.

R Core Team, 2021. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria, URL: https://www.R-project.org/.

Raffin, A., Hill, A., Gleave, A., Kanervisto, A., Ernestus, M., Dormann, N., 2021. Stable-Baselines3: Reliable reinforcement learning implementations. J. Mach. Learn. Res. 22 (268), 1–8, URL: http://jmlr.org/papers/v22/20-1364.html.

Rodríguez, A., Cordera, R., Alonso, B., dell'Olio, L., Benavente, J., 2022. Microsimulation parking choice and search model to assess dynamic pricing scenarios. Transp. Res. A: Policy Pr. 156, 253–269. http://dx.doi.org/10.1016/j.tra.2021.12.013, URL: https://linkinghub.elsevier.com/retrieve/pii/S0965856421003281.

Rodríguez, A., dell'Olio, L., Moura, J.L., Alonso, B., Cordera, R., 2023. Modelling parking choice behaviour considering alternative availability and systematic and random variations in user tastes. Sustainability 15 (11), 8618. http://dx.doi.org/10.3390/su15118618, URL: https://www.mdpi.com/2071-1050/15/11/8618.

Roßmann, J., 2010. Data quality in web surveys of the German longitudinal election study 2009. In: 3rd ECPR Graduate Conference, Vol. 30.

Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O., 2017. Proximal policy optimization algorithms. arXiv e-print URL: http://arxiv.org/abs/1707.06347.

Shoup, D., 2011. The High Cost of Free Parking, Updated Planners Press, American Planning Association, Chicago, OCLC: ocn742069668.

Shoup, D., 2021. Pricing curb parking. Transp. Res. A: Policy Pr. 154, 399–412. http://dx.doi.org/10.1016/j.tra.2021.04.012, URL: https://linkinghub.elsevier.com/retrieve/pii/S0965856421001105.

Sutton, R.S., Barto, A.G., 2018. Reinforcement Learning: An Introduction, second ed. In: Adaptive computation and machine learning, MIT Press, Cambridge, Mass.

Tchervenkov, C., 2022. Empirical and Simulation Studies on Parking in Switzerland (Ph.D. thesis). http://dx.doi.org/10.3929/ETHZ-B-000593365, URL: http://hdl.handle.net/20.500.11850/593365. Artwork Size: 202 p. Medium: application/pdf Pages: 202.

Thompson, R.G., Richardson, A.J., 1998. A parking search model. Transp. Res. A: Policy Pr. 32 (3), 159–170. http://dx.doi.org/10.1016/S0965-8564(97)00005-0, URL: https://linkinghub.elsevier.com/retrieve/pii/S0965856497000050.

Traets, F., Sanchez, D.G., Vandebroek, M., 2020. Generating optimal designs for discrete choice experiments in *R* : The **idefix** package. J. Stat. Softw. 96 (3), http://dx.doi.org/10.18637/jss.v096.i03, URL: http://www.jstatsoft.org/v96/i03/.

Train, K., 2009. Discrete Choice Methods with Simulation, second ed. Cambridge University Press, Cambridge, OCLC: 839026301.

van Nieuwkoop, R.H., Axhausen, K.W., Rutherford, T.F., 2016. A traffic equilibrium model with paid-parking search. p. 40. http://dx.doi.org/10.3929/ETHZ-A-010609175, URL: http://hdl.handle.net/20.500.11850/155508. Artwork Size: 40 p. Medium: application/pdf Publisher: [object Object].

Waraich, R.A., Axhausen, K.W., 2012. Agent-based parking choice model. Transp. Res. Rec. 2319 (1), 39–46. http://dx.doi.org/10.3141/2319-05.

Weis, C., Vrtic, M., Widmer, P., Axhausen, K.W., 2011. Influence of parking on location and mode choice: A stated choice survey. p. 18. http://dx.doi.org/10.3929/ETHZ-A-006543999, URL: http://hdl.handle.net/20.500.11850/38059.

Wilensky, U., 1999. NetLogo. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL, URL: http://ccl.northwestern.edu/netlogo/.

Willsher, K., 2023. Paris to charge SUV drivers higher parking fees to tackle ‘auto-besity’. Guardian URL: https://www.theguardian.com/world/2023/jul/11/paris-charge-suv-drivers-higher-parking-fees-tackle-auto-besity.

Yan, Q., Feng, T., Timmermans, H., 2023. A model of household shared parking decisions incorporating equity-seeking household dynamics and leadership personality traits. Transp. Res. A: Policy Pr. 169, 103585. http://dx.doi.org/10.1016/j.tra.2023.103585, URL: https://linkinghub.elsevier.com/retrieve/pii/S0965856423000058.

Yan, T., Fricker, S., Tsai, S., 2020. Response burden: What is it and what predicts it? In: Beatty, P., Collins, D., Kaye, L., Padilla, J.L., Willis, G., Wilmot, A. (Eds.), Advances in Questionnaire Design, Development, Evaluation and Testing, first ed. Wiley, pp. 193–212. http://dx.doi.org/10.1002/9781119263685.ch8, URL: https://onlinelibrary.wiley.com/doi/10.1002/9781119263685.ch8.

Zhang, K., Batterman, S., Dion, F., 2011. Vehicle emissions in congestion: Comparison of work zone, rush hour and free-flow conditions. Atmos. Environ. 45 (11), 1929–1939. http://dx.doi.org/10.1016/j.atmosenv.2011.01.030, URL: https://linkinghub.elsevier.com/retrieve/pii/S1352231011000586.