

Article

A Three-Level Framework for Spatial Demand Estimation and Latent-Demand Reconstruction in Paid Parking Systems

Ivan Eremin ¹, Petr Nikitin ^{1,*}, Rimma Gorokhova ² and Natalia Bespalova ³

¹ Department of Artificial Intelligence, Faculty of Information Technology and Big Data Analysis, Financial University Under the Government of the Russian Federation, Moscow 125167, Russia; ivan-erx@yandex.ru

² Department of Information Technologies, Faculty of Information Technology and Big Data Analysis, Financial University Under the Government of the Russian Federation, Moscow 125167, Russia; rigorokhova@fa.ru

³ Department of Mathematics and Data Analysis, Faculty of Information Technology and Big Data Analysis, Financial University Under the Government of the Russian Federation, Moscow 125167, Russia; nvbespalova@fa.ru

* Correspondence: pvnikitin@fa.ru

Highlights

- A unified three-level framework links aggregate demand dynamics, latent-demand reconstruction under capacity constraints, and spatial demand estimation over the road network.
- Censored demand at fully occupied parking zones is reconstructed from a beta arrival profile and cohort-specific parking-duration survival curves.
- The framework supports scenario analysis of tariff changes, capacity changes, and the expected utilisation of newly introduced parking zones.
- Interpretable regression results show that short stays are driven mainly by centrality (and, directionally, commercial activity), whereas longer stays are more price-sensitive.

Abstract

Estimating paid parking demand is difficult when transaction data are censored by full occupancy and when tariff changes create structural breaks. This study proposes an interpretable three-level framework combining (i) dynamic modelling of the cumulative client base and aggregate parking demand, (ii) recovery of model-implied latent potential demand in capacity-constrained zones using a transferred beta arrival profile and cohort-based survival curves, and (iii) spatial allocation of capacity-adjusted demand using road-network accessibility, land use, tariffs, capacity, and competitive-supply factors. Demand is estimated separately for four parking-duration groups using transaction data from Vladivostok, Russian Federation. Three validation questions are kept distinct. A one-origin, 15-month aggregate revenue scenario forecast after the June 2025 tariff change achieved MAPE 6.5% versus 18.8% for a seasonal-naive benchmark. Leave-one-zone-out spatial cross-validation produced OLS R^2 values of 0.49–0.63 across duration groups (0.61 for total sessions), whereas chronological cross-regime transfer of the zone model yielded R^2 values near zero or below. The evidence therefore supports spatial demand estimation, latent-demand reconstruction, and transparent scenario analysis; it does not support a claim of robust short-term temporal forecasting for zone-level demand.

Academic Editor: Theophilos Papadimitriou

Received: 28 August 2026

Revised: 2 October 2026

Accepted: 15 October 2026

Published: 18 October 2026

Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: spatial parking-demand estimation; latent demand; censored demand; road network; tariff scenarios; interpretable modelling; out-of-sample validation

1. Introduction

Parking demand in cities has been extensively studied in transport economics, urban planning, and intelligent transport systems. Accurate estimation of parking demand is essential for managing urban congestion and optimising parking infrastructure allocation [1].

One of the most influential conceptual models in parking research was proposed by Donald Shoup. In his seminal book *The High Cost of Free Parking*, Shoup demonstrated that parking-pricing policies strongly influence travel behaviour and urban congestion, arguing that inefficient parking management can significantly distort urban transport systems [2]. An empirical confirmation of these ideas was the SFpark project in San Francisco, where parking prices were adjusted according to actual occupancy: it was shown that demand responds statistically significantly to tariff changes, and the estimated price elasticity of parking demand is negative [3,4]. The issues of optimal load-driven pricing under demand uncertainty are developed in the works on stochastic tariff management [5].

Analytical models of parking demand have been further developed within transport economics. For example, the integrated model proposed by Richard Arnott and Eren Inci describes the equilibrium relationship between parking supply, congestion, and pricing mechanisms [6]. These models provide important theoretical insights into the interplay between parking capacity and urban traffic but are typically formulated at an aggregate level.

With the growing availability of digital transaction and sensor data, parking and traffic-demand studies increasingly employ empirical and data-driven methods. Recent forecasting research combines classical time-series methods with deep learning [7], while city-wide parking models increasingly exploit spatial dependencies and graph learning, including prediction under partial sensor coverage and more recent graph-Transformer formulations [8–10]. Statistical regression and machine-learning approaches remain particularly useful when the objective includes interpretable relationships between demand and explanatory variables such as price, land-use characteristics, and proximity to commercial centres [11,12].

Real-time parking occupancy and vacancy modelling has progressed from sensor-based prediction and deep learning [13–18] toward graph-based city-wide vacancy prediction and direct parking recommendation [8–10,19]. These developments improve operational prediction and guidance, but their objectives differ from estimating the latent demand that is suppressed when capacity is binding.

Spatial modelling is another important direction in parking-demand research. Many studies employ gravity-type or spatial-interaction models, where demand is distributed across zones according to attractiveness and distance from demand sources [6,20,21]. Recent work in forecasting has also revisited synthetic demand-flow generation and highlighted robustness issues of conventional gravity formulations under difficult distance configurations [22]. In parking applications, Euclidean distance is still frequently used as a proxy for accessibility, even though the actual urban road network may contain barriers and strongly asymmetric connections.

Recent research has also expanded spatio-temporal demand modelling using large-scale mobility data, synthetic activity plans, and graph representations [8,16,17,23,24]. These methods demonstrate the value of combining temporal structure with spatial

context, but they are primarily designed for prediction or activity generation rather than for correcting demand observations censored by parking capacity.

A rapidly developing research direction focuses on short-term demand and vacancy forecasting using machine-learning and deep-learning methods. Recent models can capture dynamic spatial dependencies and missing-sensor patterns with high predictive accuracy [7,8], and graph-based ranking can move directly from prediction to parking recommendation [19]. For municipal tariff and capacity planning, however, predictive accuracy alone is insufficient when the model does not explicitly distinguish observed transactions from capacity-suppressed demand or expose the structural contribution of pricing, accessibility, and competing supply.

A critical limitation of many empirical studies is the implicit assumption that observed parking sessions fully reflect underlying demand. In practice, this assumption can be violated when parking areas operate at full capacity: additional potential arrivals may be unable to initiate a recorded session, so transaction activity becomes a censored representation of preferred-zone demand. The problem of censored observations has been well studied in econometrics [25] and revenue management [26]. In transportation, latent-demand estimation has been addressed using censored Gaussian processes and censored quantile regression for shared mobility [27,28]. For parking systems, however, capacity-suppressed demand is still considered only in a limited number of studies; treating fully occupied zones as ordinary observations can systematically understate demand in the most attractive locations.

The 2023–2026 literature therefore shows substantial progress in traffic forecasting [7], synthetic mobility and demand generation [22,24], city-wide and facility-level parking-occupancy prediction [8–10], graph-based parking recommendation [19], and parking-sharing or supply-and-pricing optimisation [29,30]. These strands are complementary, but they generally address forecasting, recommendation, pricing, or spatial interaction as separate analytical tasks.

Related Russian research examines regional heterogeneity in transport-service demand using cluster analysis [31] and assesses the structural stability of transport and logistics networks using graph models [32]. These regional and network-level studies motivate attention to spatial heterogeneity and connectivity; they do not directly validate the parking-demand reconstruction proposed here.

Research involving Financial University authors also links transport forecasts to intermodal freight planning [33], examines operating costs relevant to public-transport tariffs [34], develops ontology-based transportation management systems [35], and studies sustainable transport-infrastructure financing [36]. These contributions provide a broader planning and information-system context, while the present study addresses capacity-censored parking transactions at the zone level.

The intended decision-support role is consistent with research on data-driven public administration [37] and intelligent decision-support architectures [38]. In this setting, interpretable estimates should expose their assumptions and validation limits rather than substitute model outputs for evidence about causal policy effects.

Against this background, three methodological gaps are especially relevant. First, transaction-based parking models often use recorded sessions as a direct proxy for demand even when full occupancy censors potential arrivals. Second, spatial models may rely on simplified distance measures rather than accessibility on the urban road network [22,39]. Third, aggregate demand generation, correction for capacity censoring, and spatial allocation are rarely connected in one interpretable estimation chain. Behavioural research further shows that the onset of parking search is itself endogenous to parking conditions [40], reinforcing the need to distinguish recorded parking events from underlying potential demand.

To address these issues, this study proposes a three-level model for estimating parking demand in cities. This structure integrates:

1. A dynamic model describing the temporal evolution of the cumulative client-base variable and overall parking demand;
2. A latent-demand recovery model that reconstructs model-implied potential demand under parking capacity constraints;
3. A spatial distribution model based on urban road-network accessibility.

The novelty of the proposed approach is fourfold. First, it explicitly treats full occupancy as a censoring mechanism and reconstructs potential demand before spatial-model estimation rather than learning directly from truncated transaction counts. Second, it links aggregate service-demand dynamics to zone-level allocation so that system-wide forecasts and local demand estimates remain internally consistent. Third, spatial accessibility and competitive supply are evaluated on the road network rather than solely through Euclidean proximity. Fourth, the model separates four parking-duration groups, allowing tariff sensitivity and urban-attractiveness effects to differ across short- and long-stay users. Thus, relative to recent vacancy-prediction, recommendation, and parking-pricing studies [8,19,29,30], the proposed framework is positioned as an interpretable demand-reconstruction and scenario-analysis layer rather than as a black-box real-time prediction system.

2. Materials and Methods

2.1. Data Description and Study Area

The empirical analysis in this study is based on transaction data obtained from a paid parking client service operating in the city of Vladivostok (Russian Federation). The dataset contains detailed information on individual parking sessions recorded in the city's parking management system.

Vladivostok is a major city in the Russian Far East, characterised by a dense road network and a high concentration of commercial and administrative buildings in the central districts, complex terrain, and the presence of a bay, which makes standard geo-analysis methods poorly applicable. As in many growing cities, parking demand in the central area significantly exceeds supply during peak hours.

The paid parking system analysed in this study covers several central urban zones where on-street parking is regulated via a digital parking management platform. Drivers can initiate and terminate parking sessions using a mobile application, web service, or SMS-based payment mechanism, allowing the system to record precise start and end timestamps [41].

This digital infrastructure generates high-accuracy behavioural data describing parking demand and user activity, enabling detailed statistical and econometric analysis of parking system dynamics.

Data sources. The dataset consists of historical parking transaction records collected by the parking service operator. Each record corresponds to a completed parking session and includes several attributes describing the parking event:

1. Parking zone identifier, describing the spatial location of the parking space;
2. Parked vehicle category;
3. Date and time of parking session start;
4. Date and time of parking session end;
5. Parking duration, calculated from the recorded timestamps;
6. Pseudonymous client identifier (used to construct client cohorts and the cumulative client base);
7. Payment amount.

These variables allow reconstruction of temporal demand patterns, spatial distribution of parking activity, and behavioural characteristics of parking users.

The study area covers 40 km² of central Vladivostok, comprising 46 paid on-street parking zones with a total of 3908 marked spaces. The analysis uses 2.37 million anonymised payment transactions recorded between January 2023 and February 2026. Figure 1 summarises the data flow across the three modelling levels; the geographical scope is defined by the study-setting description above.




Figure 1. Data flow across the three modelling levels: aggregate demand and revenue (Level 1), latent-demand recovery in capacity-constrained zones (Level 2), spatial allocation over the road network (Level 3), and the tariff-scenario integration step.

2.2. Mathematical Formulation of the Three-Level Demand Model

2.2.1. Overall Model Structure

Urban parking demand is determined by a combination of temporal dynamics, spatial distribution of trips, and infrastructure capacity constraints. To account for these mechanisms, a three-level model is proposed:

1. Level 1 models the temporal dynamics of aggregate parking demand and the cumulative client-base variable.
2. Level 2 reconstructs latent parking demand in zones where observed demand is constrained by capacity limitations.
3. Level 3 estimates the spatial distribution of demand across parking zones using adjusted demand values.

By reconstructing latent demand before training the spatial model, the proposed structure reduces bias and improves parameter consistency.

Let: t —time (calendar months); $CCB(t)$ —cumulative client-base variable (volume at time t); $Q(t)$ —aggregate parking demand, i.e., total number of parking sessions (volume over period t); and $Q_i(t)$ —demand for parking zone i .

The overall model can be expressed as follows:

$$Q(t) \rightarrow \tilde{Q}_i(t) \rightarrow Q_i(t), \quad (1)$$

where

$Q(t)$ —aggregate demand generated by the cumulative client-base variable;
 $\tilde{Q}_i(t)$ —capacity-adjusted demand;
 $Q_i(t)$ —spatially distributed demand.

2.2.2. Level 1 Formulation of Aggregate Demand Dynamics

The first level describes the trajectory of the client base. In the available transaction data, a client enters CCB once they have performed at least one parking session from the project launch up to month t . Thus, the variable is operationally a cumulative ever-active client base rather than a rolling active-user measure; this distinction is retained in the aggregate model.

The discrete client-base dynamics are represented by the multiplicative growth factor $r(t)$:

$$CCB_t = CCB_{t-1} \cdot r_t, \quad (2)$$

where $r(t) = CCB(t)/CCB(t-1)$ is the period-over-period multiplicative change factor.

For a continuous-time approximation, let $g(t)$ denote the instantaneous relative growth rate of the client base. The resulting dynamics are written as follows:

$$\frac{dCCB}{dt} = CCB(t) g(t), \quad (3)$$

For monthly observations, the empirical relative growth rate is estimated on the logarithmic scale to remain consistent with multiplicative growth:

$$g_t = \frac{\ln CCB_t - \ln CCB_{t-1}}{\Delta t} = \frac{\ln r_t}{\Delta t}, \quad (4)$$

where $\Delta t = 1$ month. This definition is preferable to the approximation $r(t) - 1$ when early month-to-month changes are large.

The empirical growth rate decreases over time and is approximated by the power function:

$$g(t) = a \cdot t^b, b < 0, \quad (5)$$

where

a —growth scale coefficient;

b —growth rate decay parameter.

The power-law form captures the empirically observed deceleration in client-base growth as the service matures. Because the fitted decay exponents are below -1 , the continuous solution cannot be anchored at $t = 0$. Instead, let $t_0 > 0$ denote the first observed period. Integrating from t_0 to t gives the following, for $b \neq -1$:

$$CCB(t) = CCB(t_0) \cdot \exp \left[\frac{a}{b+1} (t^{b+1} - t_0^{b+1}) \right], b \neq -1 \quad (6)$$

For the limiting case $b = -1$, the corresponding solution is as follows:

$$CCB(t) = CCB(t_0) \left(\frac{t}{t_0} \right)^a \quad (7)$$

The total number of parking sessions generated by the user base is modelled using a nonlinear regression function:

$$Q_t = k \cdot CCB_t^\gamma S_{m(t)}, \quad (8)$$

where S_m is the seasonal factor for calendar month m . To preserve the annual level, the factors are normalised as follows:

$$\frac{1}{12} \sum_{m=1}^{12} S_m = 1, \quad (9)$$

where

k —scaling parameter;

γ —reflects the elasticity of parking usage with respect to the user base. When $\gamma < 1$, session growth slows as the user base increases, while $\gamma > 1$ reflects network effects leading to higher service usage.

Combining the demand relation with the average fee per parking session yields the revenue equation:

$$R(t) = p \cdot k \cdot CCB(t)^\gamma + c, \quad (10)$$

where p is the average fee per parking session and c is an empirically estimated intercept. Equation (10) is a within-regime relation, not a tariff-invariant causal equation. For estimation, the pre-change coefficients are fitted only to observations available through May 2025. After the June 2025 tariff change, the pre-change forecast is multiplied by a fixed scenario factor derived from Level 3 predictions under the old and new tariff configurations; the post-change actual revenue is reserved for evaluation and is not used to estimate that factor.

Model assumptions: the proposed forecasting model is developed on the basis of several explicit assumptions that define its applicability domain and ensure interpretability of the results:

1. Stability of the institutional and operating environment over the forecasting horizon. Major regulatory reforms, abrupt tariff changes not represented in the scenario, competitor entry, and prolonged service disruptions are treated as exogenous structural breaks rather than endogenous model components.
2. Structural stability of the aggregate demand–revenue relationship within a tariff regime. Seasonal effects are represented empirically; persistent changes in customer behaviour, urban policy, or tariffs require model re-estimation. In particular, the tariff change introduced in June 2025 is treated as a structural break rather than as evidence that a single revenue-per-session coefficient is invariant over the full sample.
3. Homogeneity of aggregate user behaviour. User behaviour is assumed to be homogeneous at the aggregate level. Individual-level heterogeneity and micro-behavioural differences are aggregated into macro-level indicators, which is suitable for revenue forecasting but limits the method's applicability for individual demand modelling.
4. Stationarity of parking session duration distribution. Short-term fluctuations are allowed, but systematic long-term changes in the mean or variance of session durations are not considered.

Under assumptions 1–4, the aggregate component provides a transparent and interpretable forecasting framework. Its extrapolative use is most defensible for recently launched services and within periods in which the tariff regime and operating environment remain sufficiently stable.

2.2.3. Level 2 Formulation of Capacity Censoring and Latent-Demand Recovery

Observed parking activity may be distorted when a parking zone reaches full simultaneous occupancy. Capacity constrains the number of vehicles occupying the zone at a given time, not the cumulative number of sessions over a day or quarter; therefore, the censoring mechanism must be expressed through the occupancy balance.

Let $O(i,j)$ denote occupancy in zone i at the end of interval j , $A_{obs}(i,j)$ the admitted and recorded arrivals during the interval, $D(i,j)$ the departures, and $C(i)$ the physical number of parking spaces. The occupancy balance is as follows:

$$O_{i,j} = O_{i,j-1} + A_{i,j}^{obs} - D_{i,j} \quad (11)$$

Under a no-queue approximation, the admitted arrivals are related to latent potential arrivals by the following:

$$A_{i,j}^{obs} = \min(A_{i,j}^{latent}, C_i - O_{i,j-1} + D_{i,j}), \quad (12)$$

where $A_{i,j}^{latent}$ is potential arrivals in the absence of capacity constraint.

Censoring occurs when occupancy is at or near the physical capacity $C(i)$ and latent arrivals exceed the spaces released by departures. Consequently, the unobserved component is associated with rejected or displaced arrivals, while capacity does not cap the total number of sessions over the observation period.

To estimate latent demand, the vehicle arrival process is considered as a continuous stochastic process describing the distribution of parking start times within a given observation interval. Let the observation interval be normalised to the unit segment:

$$t \in [0, 1], \quad (13)$$

where t is the normalised time within the considered period (e.g., a day or working day).

It is assumed that the probability of a vehicle arriving at time t is described by a probability density function. A beta distribution is used to approximate the empirical distribution. A similar pattern of distribution of the influx of cars into the parking lot was noted earlier [42]. The density function is as follows:

$$f(t) = \frac{t^{\alpha-1}(1-t)^{\beta-1}}{B(\alpha, \beta)}, \quad (14)$$

where t —normalised observation time;

$\alpha > 0, \beta > 0$ —shape parameters;

$B(\alpha, \beta)$ —beta function serving as a normalisation constant.

The beta function is defined as follows:

$$B(\alpha, \beta) = \int_0^1 t^{\alpha-1} (1-t)^{\beta-1} dt \quad (15)$$

For quarter-hour or hourly bins, the beta density is converted to probability mass by integrating over each interval rather than by summing raw PDF ordinates:

$$p_j = F_{Beta}(t_{j+1}) - F_{Beta}(t_j) \quad (16)$$

The interval probabilities are normalised so that:

$$\sum_j p_j = 1 \quad (17)$$

Parking sessions are then grouped into hourly cohorts (generations), representing vehicles arriving within the same time interval. For each cohort k (k corresponds to the interval $[t_k, t_k + \Delta t]$ and groups all vehicles that started a parking session within that interval), the parking duration of a vehicle in cohort k is denoted D_k , with survival function

$$S_k(u) = P(D_k > u), \quad (18)$$

estimated empirically from completed sessions. The number of cohort- k vehicles still parked at time t is $N_k(t) = A_k \cdot S_k(t - t_k)$, where A_k is the number of arrivals in cohort k and t_k the cohort arrival time.

Total parking occupancy at time t is then calculated as a superposition of active cohorts:

$$O(t) = \sum_k N_k(t), \quad (19)$$

where $N_k(t)$ is the number of vehicles from cohort k still occupying parking spaces at time t .

Using this formulation, latent potential arrivals can be reconstructed for intervals in which the observed admitted flow is constrained by occupancy. The resulting quantity is a model-implied preferred-zone demand under the fitted arrival process; it is not a directly observed count of rejected vehicles.

Model assumptions: the latent-demand recovery model is based on the following assumptions:

1. Vehicle arrivals are represented by a stochastic arrival process whose intensity varies over the observation interval.
2. No explicit behavioural response to full occupancy is modelled. Drivers may wait, relocate to another zone, abandon the trip, or shift the start of their parking search; recent empirical work confirms that the onset of cruising for parking is behaviourally endogenous [40]. Because these actions are not observed in the transaction records, the reconstructed latent demand should be interpreted as potential demand under the fitted arrival process rather than as a directly observed count of rejected vehicles.
3. Invariance of session duration distribution within a cohort. Cohort survival curves are estimated from completed sessions, so it is assumed that the behaviour of clients within a cohort is equivalent for those who successfully parked and those who left.
4. Identifiability of latent demand from the uncensored tails. The shape parameters (α , β) of the intraday arrival profile are not estimated separately for each censored zone. Instead, a single common shape is estimated as the median intraday arrival profile across the zones that do not reach full occupancy (donor zones), and for each censored zone, only a scale factor is fitted on the uncensored intervals U (the pre-saturation morning limb and the post-peak evening limb). Transferring the shape from unsaturated zones stabilises identification, because a censored zone provides too few uncensored points to determine α , β and the scale simultaneously; the censored mid-day peak is then recovered by extrapolation. The procedure assumes that the uncensored tails are free of censoring, that the common intraday shape applies to the censored zones, and that the beta family adequately describes the profile—assumptions examined empirically in Sections 2.5 and 3.2.

2.2.4. Level 3 Formulation of Spatial Demand and Network Accessibility

The third level addresses the problem of modelling the spatial distribution of parking session demand in the urban environment. Unlike the aggregate level, which analyses the dynamics of the cumulative client-base variable, this stage focuses on estimating the intensity of parking space usage at the local zone level.

Instead of Euclidean distances, we use shortest-path metrics on the transport network, which allows proper accounting for barriers and the road-network structure [43].

Let the urban territory be discretised into a set of zones ($i = 1, \dots, N$). For each zone, data on the number of parking sessions are observed.

The goal is to build a demand function:

$$Q_i(t) = f(X_i), \quad (20)$$

where X_i is a vector of observed zone characteristics.

To properly describe parking session demand, the spatial structure of the urban environment must be considered. The study uses factors reflecting each zone's attractiveness and functional saturation, calculated from geographic information data.

The urban transport system is represented as a weighted graph:

$$G = (V, E), \quad (21)$$

where

V —set of nodes (intersections);

E —set of edges representing road segments.

Each edge $e \in E$ is associated with a travel cost w_e , representing distance or travel time.

For each parking zone i , the distance to demand sources is defined as the shortest-path distance in the road network:

$$d_i = \min_{p \in P} \sum_{e \in p} w_e, \quad (22)$$

where P denotes the set of feasible paths.

To account for local spatial structure, a zone attractiveness indicator is introduced, based on the density of points of interest (POIs) available within a given pedestrian accessibility radius, consistent with gravity and spatial-interaction principles [22,44]. Distances to objects are calculated over the road network, ensuring that actual connectivity and urban barriers are represented explicitly.

The integral attractiveness of a zone is defined as the sum of contributions from individual objects:

$$I_{\text{total}} = \sum_g I_g, \quad (23)$$

Here I_g denotes the attractiveness contribution of object g , and I_{total} is the sum of those contributions over all eligible objects. Capital italic I denotes an attractiveness index, not a logarithm or the letter l . Each contribution combines a point-of-interest weight with distance decay measured along the pedestrian-accessible network; the sum is therefore a network-weighted exposure index rather than a direct business count.

To account for the influence of neighbouring paid zone supply on demand, a supply pressure index was calculated over paid neighbouring zones within the parking service only, reflecting the density of competitive priced supply weighted by spatial proximity and price:

$$SPI_i = \sum_{j \neq i} \frac{Capacity_j}{Price_j \cdot Dist_j}, \quad (24)$$

where SPI_i —supply pressure indicator for zone i ;

$Capacity_j$ —capacity (number of spaces) in zone j ;

$Price_j$ —hourly tariff in zone j ;

$Dist_j$ —the road-network distance between zone centroids.

Equation (24) is evaluated only for paid neighbouring zones with $Price_j > 0$; the distance is floored at $Dist_{\text{min}} = 25$ m to avoid singularities. Free parking outside the service is represented separately by $free_places_200m$ and is therefore never introduced through a division by a zero tariff.

The use of network distances and aggregated attractiveness indicators yields spatially consistent factors suitable for use in demand models and consistent with gravity model assumptions.

problem rather than a road connection. Russian text belongs to the geographic basemap; “улица” means “Street”. Map data © OpenStreetMap contributors, ODbL.

To account for heterogeneity of the surrounding built environment, the centreline is discretised with a fixed step h (in metres) (Figure 3). At each discretisation point, a local buffer of radius R is formed, describing straight-line distance in metres to each object on the map obtained from OpenStreetMap. A pedestrian route to each object is built using the graph, then objects with pedestrian distance less than $\tilde{R}$ metres are aggregated.




Figure 3. Spatial sampling along parking zone 320 on Pushkinskaya Street. Black dots are sampling locations at 100 m intervals, the red line is the ordered zone centreline, and blue circles are 200 m Euclidean candidate-search buffers. Their overlap is intentional. Final accessibility is evaluated by pedestrian-network distance, not by circle membership alone. Russian labels are geographic base-map names; “улица” means “Street”. Map data © OpenStreetMap contributors, ODbL.

Free parking places from OpenStreetMap are aggregated in a similar manner. The final set of explanatory factors is as follows:

1. cafe_cnt;
2. government_cnt;
3. health_cnt;
4. office_cnt;
5. restaurant_cnt;
6. retail_cnt;
7. supply_pressure_index;
8. free_places_200m;
9. POI;
10. Capacity;
11. Price;
12. Walking time to centre (min).

The computational analysis was performed using freely available software in the Anaconda distribution with Python 3.6.

To account for potential nonlinearities, a regression tree model [46] (Decision-TreeRegressor, scikit-learn [47]) was additionally employed as a nonlinear reference ($\text{max_depth} = 5$, $\text{min_samples_leaf} = 4$, $\text{random_state} = 42$). Given the small training sample ($n = 41$), the tree is included only as a coarse nonlinear benchmark rather than a competitive model: with these settings, it tends to overfit the training data (high training R^2) while transferring weakly, and it is reported for comparison, not recommended for deployment.

Unlike linear regression, decision trees can capture nonlinear relationships and factor interactions without prior specification of the functional form and, most importantly, are well suited for interpretation.

To improve linear model stability under multicollinearity, ridge regularised regression was additionally applied [48]. Poisson regressions and the negative-binomial model, which are standard for modelling count data [49], are considered as computational specifications.

On the reduced core specification, all variance inflation factors are below 3 (maximum 2.8) and the condition number is 5.9×10^3 ; the elevated condition number reflects differences in variable scale rather than harmful collinearity, which is confined to the redundant free parking counts excluded from the core model (VIF up to 1.7×10^4 in the full set).

To select the most accurate model, several specifications and interpretable model variants were considered:

1. Multiple linear regression;
2. LOG-model (linear regression on $\log(Y)$);
3. Poisson (classic for count data);
4. Negative binomial;
5. Ridge;
6. Decision tree.

The revised validation protocol treats feature screening and model tuning as training-sample operations; the held-out sample is reserved for final performance evaluation. Inferential interpretation of the final OLS coefficients requires standard errors and confidence intervals in addition to point estimates.

2.4. Integrated Demand Estimation

The integrated estimate follows the three modelling levels defined above. The aggregate component supplies the system-wide demand trajectory; the occupancy balance, interval arrival probabilities and cohort survival curves provide the capacity-censoring correction; and the spatial specification allocates adjusted demand over the network. The reconstruction is performed before spatial-model estimation.

This integrated structure enables forecasting of both aggregate parking service demand and its spatial distribution, while accounting for distortions in observed data due to capacity shortages.

2.5. Model Validation and Statistical Evaluation

Model assessment distinguishes three non-equivalent questions and does not pool their metrics: (1) aggregate temporal scenario forecasting at Level 1, evaluated from a single forecast origin (31 May 2025) over horizons $h = 1, \dots, 15$ months; (2) chronological cross-regime transfer of the Level 3 zone model, trained on March–May 2025 and tested on June–August 2025; and (3) spatial transfer to an unseen zone, evaluated by leave-one-zone-out (LOZO) cross-validation. Training R^2 is reported only as goodness of fit. Temporal revenue forecasts are compared descriptively with a seasonal-naïve benchmark using MAE, RMSE, MAPE, WAPE, signed error, and scale-free MAE and RMSE skill scores; zone-level OLS models are reported with held-out R^2 , MAE, RMSE, and WAPE; count models are evaluated by Poisson deviance and deviance explained. Preprocessing, variable screening, regularisation choices, tree hyperparameters, seasonal coefficients, and prediction bounds are fixed using training data before held-out evaluation.

The observation unit is a parking zone; each zone contributes one row per target, formed by summing paid parking sessions on working days within the window and splitting them by duration bucket. The chronological hold-out uses spring 2025 (March–May 2025) for training and summer 2025 (June–August 2025) for testing; the forecast origin is the end of May 2025. Training covered 41 zones (41 rows) and the test 46 zones (46 rows), of which 40 are common and 6 appear only in the test set (zones 235, 325, 415, 425, 440,

450). Because all zones are observed for the same duration, no time-exposure offset is used for the Poisson and negative-binomial specifications. The dependent variables are absolute session counts, while capacity enters as an explanatory variable. A per-space specification is not reported and no intensity-based robustness claim is made. Poisson and negative-binomial fits are compared on the same held-out deviance scale.

Because the application concerns both forecasting over time and transfer to newly configured parking zones, the preferred validation design combines (i) a chronological hold-out, in which the training period strictly precedes the test period, and (ii) a zone-level hold-out, in which complete zones are excluded from parameter estimation. The chronological hold-out comprised $n = 41$ zone-level training observations and $n = 46$ test observations; 40 zones appear in both partitions, and 6 zones (235, 325, 415, 425, 440, 450) appear only in the test period, so the temporal hold-out combines re-observation of known zones under the new tariff regime with a small set of previously unseen zones. Because zone re-appearance makes the chronological split only a partial test of spatial transfer, generalisation to new zones is additionally assessed by leave-one-zone-out (LOZO) cross-validation, in which each zone is held out in turn and never appears simultaneously in estimation and evaluation. Feature screening used a pre-specified core specification (price, capacity, road-network distance to centre, restaurant density) rather than data-driven elimination; across targets, this core matched or exceeded a backward-elimination variant under LOZO.

The model comparisons are descriptive rather than claims of statistically significant superiority. A Diebold–Mariano test is designed for a time-ordered sequence of paired forecast errors; it is not appropriate for the cross-sectional LOZO errors reported in Section 3.3, and the three-month chronological zone hold-out is too short for reliable long-run variance estimation. Accordingly, the manuscript reports effect sizes and held-out losses, avoids pairwise significance claims, and treats a blocked or repeated temporal evaluation as necessary future work.

$$\text{Train: } t \leq t_0, \quad \text{Test: } t > t_0.$$

$$Z_{\text{train}} \cap Z_{\text{test}} = \emptyset.$$

For the latent-demand component, the reconstruction is validated by a pseudo-censoring experiment on the zones that do not reach full occupancy, for which the full arrival profile is observed. For each such zone, the midday interval (09:00–16:00, the window censored in the saturated zones) is withheld, the common intraday shape is estimated from the remaining zones, only the scale is fitted on the retained morning and evening limbs, and the reconstructed midday profile is compared with the withheld observations using the recovery ratio and the mean absolute error, root-mean-square error, and bias of the reconstructed hourly shares. The beta family is benchmarked against alternative smooth profiles (gamma, log-normal, and a model-free shape-preserving interpolation of the tails), and the sensitivity of the reconstructed totals to the hourly discretisation of the beta density is quantified. Confidence intervals for the reconstructed totals were obtained by a day-level bootstrap: whole working days were resampled with replacement ($B = 1000$), and in each replication, the city-wide arrival shape was held fixed while the zone-specific scale was re-estimated on the uncensored limbs and the total was recomputed; 95% percentile intervals are reported in Section 3.2. Whole days rather than individual hours were resampled because arrival counts in adjacent within-day intervals are dependent.

To test whether reconstruction accuracy changes with saturation severity, a graded pseudo-censoring experiment was performed on the common subset of 37 unsaturated zones that could be evaluated under every censoring condition. Four central windows were withheld: 10:00–15:00, 09:00–16:00, 08:00–16:00, and 09:00–17:00. They concealed

median shares of 60.8%, 79.3%, 85.8%, and 86.8% of observed daily arrivals, respectively. At each level, the common intraday shape was estimated without the target zone, the scale was fitted on the retained limbs, and the recovery ratio was calculated against the known withheld mass. This common-subset design separates the effect of censoring severity from changes in zone composition.

3. Results

3.1. Aggregate Demand Dynamics and Revenue Scenarios

The monthly growth dynamics of the cumulative client-base variable CCB show a clear decelerating trend (see Figure 4), from approximately 100% in the first month to 1–2% in later months. The empirical log-growth rate exhibits seasonal variation, with faster expansion in summer months.

For comparison, a parallel forecast is built using the same model but with data from the first six months of the project (the main source of inaccuracies in this forecast is the lack of second-half seasonality data).




Figure 4. Observed monthly log-growth of the cumulative client base and power-law fits. Black denotes the January 2023–December 2024 calibration (24 monthly observations); blue denotes January–June 2023 (6 observations). Solid segments show observations and dashed lines show the corresponding fitted power laws. The vertical range is restricted to 0–20% to display later-month dynamics; the initial growth rate is outside that range.

Next, the monthly number of parking sessions is related to the cumulative client base through the nonlinear elasticity and multiplicative seasonal specification defined in Section 2.2.2 (Figure 5).




Figure 5. Monthly parking sessions versus the cumulative client base. Black points represent January 2023–December 2024 (24 observations) and blue points January–June 2023 (6 observations). Dashed lines show the corresponding power-law fits.

Within a fixed tariff regime, parking revenue is approximated as a linear function of the number of sessions (see Figure 6). Because new tariff rates were introduced in June 2025, the average revenue-per-session coefficient is treated as regime-specific rather than invariant across the structural break. The pre-break calibration is shown below:

$$R(t) = 92.163 \cdot Q(t) - 572,743 \quad (30)$$




Figure 6. Monthly revenue versus parking-session counts under the pre-change tariff regime. Black points denote January 2023–December 2024 and blue points January–June 2023; dashed lines are the corresponding linear calibration relations.

Finally, a parking revenue forecast is produced (see Figure 7).




Figure 7. Actual monthly revenue (black solid line), the forecast calibrated through December 2024 (orange dashed line), and the forecast calibrated through June 2023 (green dashed line). New tariff rates were introduced in June 2025. Values are in RUB; forecasts beyond the last observed month are scenario projections.

3.2. Latent-Demand Reconstruction and Sensitivity Results

The intraday flow of recorded arrivals exhibits a smooth unimodal structure in zones that do not reach full occupancy. For Zone 40, a beta curve provides a useful empirical approximation ($R^2 = 86.4\%$; Figure 8), supporting its use as a candidate parametric family for the uncensored arrival profile rather than as an assumed universal law.




Figure 8. Intraday recorded arrival shares for zone 40 on weekdays in March–May 2025. The black line is the fitted beta profile and the blue line the observed quarter-hour arrival shares.

At the same time, zones with 100% occupancy exhibit an anomaly (see Figure 9). While occupancy reaches a maximum of 100% (see Figure 10), the inflow of vehicles drops sharply to nearly zero and remains at a level equal to the outflow, as only those users for whom a space is available (i.e., those who register and pay) enter. Zone 410 shows no such anomaly, indicating that its demand is not censored.




Figure 9. Intraday recorded flow shares on weekdays in March–May 2025: (a) arrivals; (b) departures. Black, blue and orange identify zones 330, 340 and 410, respectively. Separate panels retain a common horizontal time scale.




Figure 10. Weekday occupancy divided by marked capacity, March–May 2025. Black, blue and orange lines identify zones 330, 340 and 410. The green dashed line marks nominal capacity (100%); values slightly above it reflect recorded occupancy relative to marked capacity.

Estimating the beta arrival profile: in zones with sufficient spare capacity, the intraday arrival profile is approximated by a beta distribution. In zone 330, the observed inflow collapses when occupancy reaches capacity, whereas the morning and late-day tails retain the expected shape. The beta parameters are separated into shape and scale. To avoid over-parameterising the reconstruction in the saturated zones—where only the morning and evening limbs are uncensored, so too few points remain to identify shape and scale jointly—a common intraday shape (α, β) is first estimated as the median arrival profile across the unsaturated donor zones; the resulting city-wide shape is $\alpha = 2.26, \beta = 2.75$ (39 donor zones). For each censored zone, only a scale factor c is then fitted by least squares on the set U of uncensored intervals,

$$\hat{c} = \operatorname{argmin}_c \sum_{j \in U} [y_j - c f_{\text{Beta}}(t_j; \alpha, \beta)]^2,$$

and the fitted profile is scaled to the reconstructed total arrival volume. Because the shape is borrowed from unsaturated zones rather than estimated on the short censored-zone tails, the recovered peak is far less sensitive to the exact choice of tail intervals than in a per-zone three-parameter fit. The reconstructed hourly volumes are obtained by integrating the fitted beta density over each hourly bin as exact interval probabilities $p_j = F_{\text{Beta}}(t_{j+1}) - F_{\text{Beta}}(t_j)$. A cruder hourly ordinate approximation differs from the reported totals by 6.2% for zone 330 and 1.3% for zone 340 (Zone 330: 51,671 vs. 54,854; Zone 340: 32,405 vs. 31,992), confirming that the totals do not depend materially on the discretisation of the density (Figure 11).

It should be emphasised that only the scale is identified from the uncensored segments; the common shape is transferred from the donor zones: the morning rising limb before the capacity-induced collapse and the late-day declining limb after demand again falls below capacity. The censored midday interval is excluded from the fit and reconstructed by extrapolation. Consequently, the recovered peak is inferred rather than observed and is sensitive to (i) whether the selected tails are genuinely uncensored, (ii) the adequacy of the beta family, (iii) the normalisation of interval probabilities, and (iv) behavioural responses to full occupancy. Identification is therefore weaker when the uncensored tails are short.

Using the fitted profile, the model-implied reconstructed potential demand for March–May 2025 is 51,671 arrivals (95% conditional bootstrap interval [49,264; 53,823]) for zone 330 and 32,405 (95% conditional bootstrap interval [31,149; 33,476]) for zone 340, compared with 15,161 and 12,503 recorded parking sessions, respectively. The intervals condition on the donor-zone set, the transferred city-wide shape parameters (α , β), and the selected uncensored intervals U ; they quantify day-level scale uncertainty rather than full pipeline uncertainty. These values should be interpreted as scenario estimates under the fitted unconstrained-arrival process rather than as directly observed suppressed demand. The corresponding amplification relative to the recorded session counts is $51,671/15,161 \approx 3.41$ and $32,405/12,503 \approx 2.59$; equivalently, the recorded counts represent approximately 29.3% and 38.6% of the reconstructed potential demand.

Zones 330 and 340 differ in both the magnitude of recorded use and the severity/timing of the capacity constraint. Zone 330 records 15,161 sessions but has the larger reconstructed total and amplification factor (51,671; 3.41), whereas Zone 340 records 12,503 sessions and has a smaller reconstructed total and amplification factor (32,405; 2.59). The forward occupancy calculation nevertheless indicates a higher unconstrained peak relative to capacity in Zone 340 (3.5 versus 2.8). Thus, the difference is driven jointly by scale, capacity, and the observed uncensored limbs; it should not be attributed to a single behavioural or tariff cause.

The procedure was first tested by pseudo-censoring the 42 zones that never reach full occupancy, for which the true midday profile is observed. For each zone, the midday window (09:00–16:00) was withheld, the shape was taken from the remaining zones, the scale was fitted on the retained limbs, and the reconstructed midday was compared with the withheld data. Across all 42 zones, the median recovery ratio (reconstructed over true withheld midday mass) is 1.05, with a median absolute error of 0.020 and a bias of -0.004 in units of the daily arrival share; on the 36 zones well described by the common shape (profile $R^2 \geq 0.80$), the recovery ratio is 0.93 and the median absolute error 0.018. Under the same protocol, the numerical median MAE and RMSE values for the beta and gamma profiles are close; the log-normal profile and shape-preserving tail interpolation have larger numerical errors, and beta has bias closest to zero (Table 1). These are descriptive comparisons, not a formal statistical ranking. A forward consistency check confirms the mechanism: reconstructed latent arrivals reproduce the observed saturation window, while unconstrained potential occupancy reaches 2.8 and 3.5 times capacity for zones 330 and 340. Passing the arrivals through the capacity constraint admits 16,633 vehicles against 15,161 recorded sessions in zone 330 and 11,392 against 12,503 in zone 340. Perturbing cohort survival curves by $\pm 20\%$ leaves the reconstructed arrival totals unchanged and shifts only peak reconstructed occupancy by about $\pm 10\%$.

The graded experiment on the fixed 37-zone subset does not show deterioration as more of the profile is hidden. For the four censoring windows listed in Section 2.5, the median recovery ratios are 1.104, 1.084, 1.003, and 0.978, while the median hidden shares increase from 60.8% to 86.8%. The corresponding absolute median proportional deviations from complete recovery, $|\text{recovery ratio} - 1|$, are 10.4%, 8.4%, 0.3%, and 2.2%. Thus, reconstruction error does not increase monotonically with saturation severity over the tested range; the estimate remains close to complete recovery even when approximately 86% of the daily profile is concealed. This result concerns median proportional bias on the common subset and does not imply that tail errors for individual zones are invariant to saturation.

Table 1. Pseudo-censoring reconstruction accuracy by arrival-profile family (medians over 42 unsaturated zones; midday 09:00–16:00 withheld). Errors are in units of the daily arrival share; the recovery ratio is reconstructed over true withheld midday mass.

Arrival Profile	MAE	RMSE	Bias	Recovery Ratio
Beta (transferred shape)	0.020	0.025	−0.004	1.045
Gamma (transferred shape)	0.020	0.024	+0.005	0.952
Log-normal (transferred shape)	0.024	0.027	+0.014	0.860
Shape-preserving interpolation	0.031	0.034	+0.030	0.699




Figure 11. Recovery of potential arrivals. Solid black, blue and orange lines are observed inflow shares for zones 330, 340 and 410; dashed green and purple lines are the reconstructed unconstrained profiles for zones 330 and 340. Time is local time of day.

Next, reconstructed arrivals are divided into hourly cohorts (“generations”, e.g., 7:00–8:00, 8:00–9:00). Cohort sizes are obtained by aggregating the normalised beta interval probabilities within each hour and scaling them to the reconstructed total arrival volume.

For each hourly arrival cohort, empirical survival curves are estimated from completed sessions recorded in March–May 2025. Each curve gives the fraction of the cohort still parked after a specified elapsed duration. Survival analysis provides the statistical framework [50]; related parking research examines search duration [51] and parking-session duration [52], which are distinct outcomes. In the present data, the early-morning cohorts remain parked for much of the working day, whereas later cohorts depart sooner (Figure 12). Occupancy is computed by superposing the surviving cohorts.

Across the 42 unsaturated zones, the recovery ratio has median 1.05 (interquartile range 0.87–1.23; 10–90% range 0.76–1.59), the reconstruction MAE of the hourly shares has median 0.020 (10–90% range 0.011–0.056), and the bias is centred near zero (median −0.005; 10–90% range −0.053 to +0.025). The distribution (Figure 13) shows a recovery ratio centred near one but with appreciable between-zone dispersion rather than relying on a single favourable case; a small number of zones over-recover by up to about twice the withheld mass, consistent with genuine deviations from the common intraday shape.




Figure 12. Cohort survival curves for zone 330, March–May 2025. Panels group arrival hours as 06:00–10:00, 11:00–15:00 and 16:00–20:00. Within each panel, the legend identifies the arrival-hour cohort by colour and line pattern. The horizontal axis is elapsed hours after arrival and the vertical axis is the fraction remaining parked.




Figure 13. Distribution of pseudo-censoring reconstruction errors across the 42 unsaturated zones: recovery ratio (reconstructed/true withheld midday mass; dashed line at 1.0) and mean absolute error of the reconstructed hourly shares.

3.3. Spatial Model Validation and Comparison

For each specification, the target was considered in absolute session counts. Training and held-out results are reported separately in Table 2; training fit is not interpreted as out-of-sample predictive accuracy.

Model assessment separates three questions. On the later held-out period, the honest out-of-sample R^2 for OLS is near zero or negative (total sessions -0.06 ; 0–1 h -0.02 ; 2 h -0.03 ; 3–4 h -0.34 ; 5+ h -0.83 on raw predictions), indicating substantial distribution shift

after the tariff change. Under LOZO cross-validation, the model is markedly more stable, with R^2 of 0.61 (total), 0.52 (0–1 h), 0.55 (2 h), 0.63 (3–4 h) and 0.49 (5+ h). Held-out MAE/RMSE/WAPE for OLS are 2258/3060/0.54 (total), 1504/2105/0.50 (0–1 h), 425/604/0.59 (2 h), 201/270/0.69 (3–4 h) and 107/158/0.88 (5+ h). Logarithmic OLS transfers poorly even after Duan’s smearing retransformation (held-out R^2 between -3.3 and -14.0). For the count specifications, held-out Poisson deviance explained is 8%, 36%, 41% and 49% across the four duration groups, exceeding the negative-binomial values; deviance rather than R^2 is the appropriate criterion for these models. Range restrictions change the clipped statistics but not this qualitative picture, and raw and clipped R^2 are reported side by side.

Table 2. Model comparison using predictive R^2 on a common response scale: training fit, honest held-out performance (raw and clipped), and leave-one-zone-out performance across duration targets. For Poisson and negative-binomial models, held-out deviance is the primary model-specific criterion and is reported in the text.

R^2	Model Specification	Total Sessions	0–1 h	2 h	3–4 h	5+ h
R^2 (training)	ols	0.713	0.656	0.669	0.740	0.655
	ridge	0.712	0.655	0.668	0.739	0.654
	poisson	0.698	0.631	0.649	0.794	0.765
	neg_bin	0.558	0.519	0.406	0.596	0.617
	log_ols	0.311	0.333	0.130	0.360	0.449
	tree	0.799	0.731	0.782	0.750	0.573
R^2 (test)	ols	−0.056	−0.019	−0.035	−0.336	−0.832
	ridge	0.004	0.029	0.027	−0.240	−0.698
	poisson	−0.027	−0.457	0.220	0.587	0.738
	neg_bin	−1.021	−1.240	−1.236	−0.074	0.497
	log_ols	−6.516	−6.820	−7.135	−1.805	−0.407
	tree	0.413	0.256	0.257	0.102	−0.380
R^2 (test, clipped)	ols	0.011	−0.018	0.109	0.283	0.234
	ridge	0.058	0.029	0.151	0.309	0.255
	poisson	−0.027	−0.457	0.220	0.587	0.738
	tree	0.413	0.256	0.257	0.102	−0.380
LOZO R^2 (test)	ols	0.606	0.523	0.554	0.629	0.491
	ridge	0.609	0.527	0.557	0.630	0.490
	poisson	0.555	0.440	0.491	0.704	0.667
	neg_bin	0.224	0.191	−0.142	0.207	0.297
	log_ols	−0.497	−0.318	−1.032	−0.512	−0.152
	tree	0.370	0.233	0.398	0.400	0.332
Test deviance	Poisson	−	55 677	13 586	6 612	3 707
	Neg. binomial	−	68 317	20 280	7 556	4 113
Deviance explained	Poisson	−	8.3%	35.7%	41.3%	48.8%
	Neg. binomial	−	−12.5%	4.0%	33.0%	43.2%

For the count specifications, Table 2 uses the same chronological hold-out defined in Section 2.5. Test deviance is the sum of Poisson unit deviances on the 46 held-out zone observations. Deviance explained is $1 - \frac{D_{model}}{D_{null}}$, where D_{null} is the held-out deviance of an intercept-only model fitted on the training sample. These quantities are not LOZO statistics. The numerical held-out deviance is smaller for the Poisson specification than for the negative-binomial specification in every duration group, and the reported deviance explained reaches 48.8% for 5+ h. The comparison is descriptive; it is not presented as a formal statistical ranking or evidence of general superiority.

Linear regression is retained as the principal explanatory specification because it combines moderate LOZO effect sizes ($R^2 = 0.49\text{--}0.63$) with transparent coefficients. Ridge gives nearly identical LOZO losses, so the evidence does not establish statistically significant superiority of OLS over ridge; the choice is interpretive rather than a claim of lower expected loss.

The model was subsequently used to evaluate scenarios involving new streets and changes in capacity and tariffs [53], and the resulting session forecasts were propagated to the revenue adjustment shown in Figures 14 and 15.

Applying the Level 3 session estimates under the summer-2025 tariff configuration and comparing them with the Level 1 pre-change revenue forecast yields a fixed scaling factor $\varphi_R = \frac{30.95}{26.68} = 1.16$. Both numerator and denominator were calculated at the forecast origin (31 May 2025) from model predictions; no revenue observed after that date entered the calculation. The factor was fixed for all 15 horizons before validation. Because it combines the transition from aggregate to spatial modelling with the tariff configuration, it is a scenario adjustment and not an estimated causal price effect.

The tariff-adjusted revenue forecast was validated against 15 months of subsequently observed revenue (June 2025–August 2026), a period spanning the June 2025 tariff change. It tracked actual monthly revenue with a mean absolute percentage error of 6.5% (median 4.4%) and was close to unbiased, overstating cumulative revenue over the period by only 0.9% (mean signed error +2.1%). The largest single-month deviation (+35% in January 2026) coincided with an atypically low observed revenue; excluding it, the mean absolute percentage error falls to about 4.5%. Because the adjustment factor was derived entirely from model predictions under the two tariff configurations, with the observed post-change revenue used only for comparison, this is an out-of-sample validation of the integrated framework across a tariff-regime change.

The forecast series $R_{(new)(t)}$ is the pre-break linear revenue model $R_{(base)(t)}$ (Equation (30)) rescaled by the fixed tariff-adjustment factor derived from the Level 3 session forecasts under the two tariff configurations.

All 15 monthly forecasts share a single forecast origin (31 May 2025); this is a one-shot multi-step forecast and the model is not re-estimated month by month. Over June 2025–August 2026, the adjusted forecast has MAE RUB 0.601 M, RMSE RUB 0.778 M, MAPE 6.5%, WAPE 5.6%, and mean signed error RUB +0.099 M. The corresponding seasonal-naïve values are MAE RUB 2.062 M, RMSE RUB 2.405 M, MAPE 18.8%, WAPE 19.3%, and mean signed error RUB −2.044 M. Relative to that benchmark, the descriptive MAE and RMSE skill scores, defined as $1 - \text{loss_model}/\text{loss_naïve}$, are 70.9% and 67.7%, respectively. Table 3 reports the monthly values. ETS and ARIMA were not used as confirmatory benchmarks because only 29 pre-break monthly observations are available, barely exceeding two seasonal cycles; fitting and selecting seasonal time-series models on such a short series would add substantial specification uncertainty. With one forecast origin and dependent multi-horizon errors, these differences are reported as effect sizes rather than tested as statistically significant superiority.

Table 3. Monthly actual revenue, seasonal-naïve baseline, and one-shot tariff-adjusted forecast over the 15-month evaluation horizon (forecast origin: end of May 2025).

Month	h	Actual, RUB	Baseline (Seasonal-Naïve)	Adjusted Forecast	Error %
June 2025	1	10,003,564	6,524,364	9,212,288	−7.9
July 2025	2	12,067,148	8,496,583	11,342,538	−6.0
August 2025	3	11,016,078	7,907,890	10,432,210	−5.3
September 2025	4	12,046,610	8,013,820	11,037,490	−8.4

October 2025	5	11,889,478	8,857,554	11,649,222	−2.0
November 2025	6	9,602,945	8,085,789	9,664,437	+0.6
December 2025	7	10,381,510	8,310,963	11,317,203	+9.0
January 2026	8	6,095,676	5,898,702	8,233,823	+35.1
February 2026	9	9,299,185	7,861,509	9,895,362	+6.4
March 2026	10	11,291,107	8,855,170	11,578,859	+2.5
April 2026	11	12,075,008	9,764,100	11,662,486	−3.4
May 2026	12	10,401,705	8,029,609	10,656,214	+2.4
June 2026	13	11,093,825	10,003,564	11,277,044	+1.7
July 2026	14	11,930,652	12,067,148	12,455,004	+4.4
August 2026	15	11,161,260	11,016,078	11,426,151	+2.4




Figure 14. Actual revenue and revenue scenarios. Black is observed revenue; orange and green are forecasts calibrated through December 2024 and June 2023, respectively; blue is the fixed tariff-adjusted scenario established at the end of May 2025. Values are in RUB.




Figure 15. Monthly revenue versus parking-session counts. Black points show the January 2023–December 2024 calibration, blue points the January–June 2023 subset, and orange points the post-change observations. Dashed lines show the displayed linear relations; these descriptive fits are not a causal tariff-effect estimate.

Linear regression trained on old data reproduces the cross-sectional demand pattern for zones with new parameters (LOZO $R^2 \approx 0.6$), although absolute transfer across the tariff change is weaker (see Figure 16).




Figure 16. Zone-level demand-estimation errors. Green points are within $\pm 40\%$ of observed demand, red points indicate overprediction by the model, and black points indicate underprediction by the model. Coincident points represent nearby parking zones and are not an additional statistical category. The $\pm 40\%$ threshold is an operational tolerance, not a prediction interval. The basemap label “ВЛАДИВОСТОК” means Vladivostok. Map data © OpenStreetMap contributors, ODbL.

3.4. Final Model Parameters

Table 4 reports the final OLS estimates with heteroscedasticity-consistent (HC3) standard errors, p -values and 95% confidence intervals ($n = 41$). Robust inference is used because Breusch–Pagan tests indicate heteroscedasticity for the 3–4 h and 5+ h groups ($p < 0.01$), while residuals are approximately normal (Jarque–Bera $p > 0.26$). Capacity and network distance to the centre are significant in every duration group; price is significant only for the 3–4 h and 5+ h groups ($p = 0.034$ and 0.042); restaurant density is not significant in any group ($p \geq 0.08$).

Table 4. Final OLS coefficients with HC3 robust standard errors, p -values in parentheses ($n = 41$).

Variable	0–1 h	2 h	3–4 h	5+ h
Intercept	2739.34 (1143.74; $p = 0.022$)	1041.03 (330.91; 0.003)	623.07 (221.31; 0.008)	375.13 (167.20; 0.031)
Price	−12.63 (11.04; 0.260) [−35.02; 9.76]	−8.27 (4.23; 0.058) [−16.85; 0.31]	−6.66 (3.02; 0.034) [−12.77; −0.54]	−4.56 (2.17; 0.042) [−8.96; −0.17]
Capacity	27.67 (5.90; <0.001) [15.71; 39.64]	9.14 (2.57; 0.001) [3.92; 14.36]	4.53 (1.65; 0.010) [1.18; 7.89]	2.87 (1.30; 0.034) [0.24; 5.50]
center_dist	−56.66 (24.12; 0.024)	−19.87 (6.75; 0.006)	−9.96 (3.32; 0.005)	−6.02 (2.81; 0.039)

	[-105.57; -7.75]	[-33.55; -6.19]	[-16.69; -3.23]	[-11.71; -0.33]
	163.07	60.75	43.86	25.96
restaurant_cnt	(153.38; 0.295)	(48.97; 0.223)	(24.23; 0.079)	(21.87; 0.243)
	[-148.00; 474.14]	[-38.56; 160.07]	[-5.1; 92.8]	[-18.40; 70.32]

4. Discussion

The obtained regression models make it possible to assess the influence of infrastructural, pricing, and spatial factors on the distribution of demand across parking zones, as well as revealing differences in user behaviour depending on parking session duration.

Capacity is positive and significant in all four groups, with the largest coefficient for short stays (about 28 additional sessions per additional space for 0–1 h, declining to about 3 for 5+ h). Because the dependent variable is an absolute count, part of this association is mechanical and should not be interpreted as a per-space intensity effect. Network distance to the centre is negative and significant throughout, and strongest for short stays (about –57 sessions per additional minute for 0–1 h). Price is retained mainly for the 3–4 h and 5+ h groups (–6.66 and –4.56) and is significant only there, indicating an inverse conditional association between tariff and longer-stay demand rather than a causal elasticity. Restaurant density is positive but not statistically significant in any group ($p \geq 0.08$) and is reported as a directional association only. The qualitative pattern is consistent with differentiated parking services and joint supply–pricing research [29,30].

Residual spatial autocorrelation is negligible: Moran's I on the OLS residuals [54,55] is not significant for any duration group (I between –0.02 and 0.10; all $p \geq 0.09$), but this does not establish that spatial dependence is absent or that spatial controls would be unnecessary in other samples.

The results allow formulating a number of practical conclusions for parking space management tasks:

1. First, the model can be used to explore differentiated-tariff scenarios by duration group. The negative price coefficients for longer stays are conditional associations in observational data; because tariffs were not randomly assigned and omitted spatial or behavioural variables may remain, they do not identify a causal price elasticity or prove that a tariff intervention will redistribute demand.
2. Second, short-term demand is largely determined by urban activity and zone centrality. This means that, in the most commercially active areas, demand management should rely not only on tariff measures but also on capacity management, organisation of drop-off/pick-up zones, and limitation of parking duration.
3. Third, the framework can represent free parking supply as a separate spatial covariate and thereby avoids treating a zero tariff inside the price-weighted SPI. However, because free_places_200m is not included in the final core specification, the present estimates do not establish an independent effect of free parking. Such an effect should be tested in a pre-specified robustness model before drawing tariff-policy conclusions.

The positioning of these results relative to the recent literature is also important. Recent graph- and Transformer-based models can now predict on-street and park-and-ride occupancy under heterogeneous spatial and contextual conditions [8–10], while heterogeneous graph ranking can directly recommend on-street spaces [19]. Recent Transportation Research Part C studies, in turn, focus on service differentiation in parking sharing and on the joint design of parking supply and pricing [29,30]. The present framework is complementary to these approaches: its main role is to reconstruct capacity-suppressed demand and produce interpretable, duration-specific demand responses that can serve as inputs to operational prediction, pricing, or allocation models.

The latent-demand estimates require cautious interpretation. The reconstructed totals of 51,671 and 32,405 are approximately 3.41 and 2.59 times the corresponding recorded session counts. Pseudo-censoring validation on unsaturated zones shows that the reconstruction recovers the withheld midday mass with a median recovery ratio near unity and small bias, and that the beta and gamma profiles have similar median absolute errors, while the log-normal and model-free alternatives have larger numerical errors; a forward occupancy check reproduces saturation over the observed censored window, and the totals are insensitive to the assumed duration distribution and to the density discretisation ($\leq 6.2\%$). The estimates nonetheless remain sensitive to the assumption that the selected tails are uncensored and to behavioural responses when a zone is full, and substitution between neighbouring zones is not explicitly modelled: a driver unable to park in one zone may appear as observed demand in another, so the reconstructed zone-level values should not be summed mechanically as independent city-wide suppressed demand. A day-level bootstrap yields 95% confidence intervals of [49,264; 53,823] and [31,149; 33,476] for the two reconstructed totals, so the estimates should be read as model-implied scenario ranges rather than precise counts of rejected vehicles.

Despite the obtained results, the developed model has several limitations:

1. First, the model is aggregate in nature and describes demand at the zone level without individual modelling of individual user behaviour. Consequently, personal client preferences, trip purposes, and repeated behaviour patterns are not considered.
2. Second, regression models assume relative stability of demand structure over time. Substantial changes in transport infrastructure, urban environment, or user behaviour may require re-estimation of model coefficients.
3. Third, the linear functional form limits the model's ability to describe nonlinear effects and complex factor interactions. In particular, the influence of price or competitive environment may change at certain thresholds.
4. Fourth, latent-demand reconstruction is identified by extrapolation from uncensored tails. Pseudo-censoring quantifies reconstruction error and benchmarks beta against gamma, log-normal, and model-free profiles, while the day-level bootstrap intervals are conditional on the selected donor zones, transferred city-wide shape, censoring-window definition, and uncensored-tail boundaries. On the common 37-zone subset, graded pseudo-censoring from 60.8% to 86.8% hidden mass does not produce a monotonic increase in absolute median proportional bias (10.4%, 8.4%, 0.3%, and 2.2%). This finding supports robustness across the tested censoring windows but does not establish invariance of zone-level tail errors. Alternative donor sets and tail boundaries should still be re-estimated in a full-pipeline bootstrap before external operational use.
5. Fifth, the chronological hold-out reuses most training zones, and honest out-of-sample R^2 under chronological transfer is near zero, so the Level 3 component is best understood as a spatial/explanatory estimator evaluated by LOZO rather than as a temporal forecaster. The final specification is based on only $n = 41$ zone observations; even with HC3 standard errors, coefficient uncertainty is material, model comparisons have low power, and rare zone types are weakly represented. Residual spatial autocorrelation was tested and found negligible, but independent city-level replication is required.

Smaller held-out errors are observed for zones with moderate demand, whereas larger errors occur for the largest central zones and some peripheral locations. This pattern reflects urban heterogeneity and the limited number of observations for rare zone types; it does not constitute a formal ranking of model performance across zone classes. Outputs should therefore be described as estimates or scenario projections, interpreted with held-

out diagnostics, and re-estimated after major changes in tariffs, capacity, or network configuration.

Despite these limitations, the proposed approach provides a high level of interpretability and a transparent structure for scenario analysis. Its predictive and policy use should be accompanied by held-out diagnostics and uncertainty analysis, particularly when tariffs, capacity, or network configuration differ materially from the estimation sample.

5. Conclusions

This paper developed an interpretable framework for spatial parking-demand estimation and reconstruction of demand censored by capacity constraints. Aggregate revenue forecasting is retained as a supporting scenario-analysis component rather than the principal contribution.

The study analysed existing parking-demand modelling methods, including classical statistical approaches, machine-learning methods, spatio-temporal models, and agent-based simulations. It was found that a significant portion of modern solutions suffer from insufficient interpretability, high data sensitivity, and limited applicability in public and municipal governance tasks.

The framework combines an aggregate demand layer, a latent-demand-reconstruction layer, and duration-specific spatial regression over road-network and land-use factors. The duration groups expose heterogeneity between short- and long-stay users while retaining a form that can be audited by municipal analysts.

Algorithms for calculating factors characterising the following were developed:

1. Transport accessibility of parking zones;
2. Volume of free parking supply;
3. Saturation with urban infrastructure objects;
4. Competitive pressure from neighbouring zones;
5. Spatial structure of the urban environment

To obtain compact and interpretable equations, the explanatory specifications were reduced to a smaller factor set. Under the revised validation protocol, all factor screening must be conducted using the training sample only, and final statistical inference should be based on coefficient uncertainty reported for the locked specification.

Empirical testing produced heterogeneous numerical out-of-sample results across targets. Leave-one-zone-out R^2 ranges from 0.49 to 0.63 across duration groups (0.61 for total sessions), whereas out-of-sample R^2 under chronological transfer is close to zero. This contrast supports interpreting the Level 3 component as a spatial/explanatory estimator rather than a temporal forecaster; it is not a formal statistical ranking of competing models. The coefficient patterns suggest systematic associations of demand with tariffs, capacity, centrality, and the competitive environment, but these relationships should be interpreted as conditional associations rather than causal effects.

The sensitivity checks have direct practical implications. Reconstructed arrival totals are stable to a $\pm 20\%$ perturbation of the cohort survival curves, while peak reconstructed occupancy changes by about $\pm 10\%$; changing density discretisation changes reconstructed totals by no more than 6.2%. In the graded pseudo-censoring experiment, increasing the median hidden share from 60.8% to 86.8% does not cause a monotonic increase in absolute median proportional bias, which takes values of 10.4%, 8.4%, 0.3%, and 2.2% across the four windows. By contrast, the transferred arrival-profile family, donor-zone composition, the boundaries assumed to be uncensored, and unmodelled substitution when a zone is full remain the more consequential or incompletely quantified assumptions. Operational use should therefore prioritise local validation of profile transfer and tail selection,

and treat duration-distribution and numerical-discretisation choices as secondary sources of uncertainty.

The practical significance of the work lies in using the software-analytical framework for transparent, assumption-dependent scenario analysis, including:

1. Exploring model-based tariff scenarios without interpreting observational price coefficients causally;
2. Estimating the likely utilisation of newly configured parking zones;
3. Analysing competitive interactions between territories;
4. Supporting decisions by city authorities and parking operators.

The framework has been demonstrated only for Vladivostok, with 46 zones and two capacity-constrained reconstruction cases. Its coefficients, arrival shape, tariff response, and validation results should not be generalised to other cities or transport services without local re-estimation and external validation. The most valuable next dataset would combine harmonised transaction records from several cities with quarter-hour occupancy or sensor counts, zone capacity and tariff histories, parking-duration distributions, road-network accessibility, points of interest, and nearby free or private supply over at least 24–36 months spanning stable and changed tariff regimes. A particularly informative validation design would leave one city out, repeat reconstruction at pre-specified saturation levels, and evaluate temporal transfer before and after tariff changes. Transferability is a hypothesis for future multi-city study, not an empirical conclusion of the present analysis.

Promising directions for further research include:

1. Finer-grained intraday, weekly, and seasonal demand dynamics;
2. Dynamic graph-based representations of urban mobility and accessibility;
3. Implementation of adaptive dynamic pricing mechanisms;
4. Expansion of behavioural factor sets;
5. Combination of interpretable models with modern machine-learning methods.

Author Contributions: Conceptualisation, I.E. and P.N.; methodology, I.E., P.N. and N.B.; software, I.E. and P.N.; validation, R.G. and N.B.; formal analysis, I.E., P.N. and N.B.; investigation, I.E. and P.N.; data curation, I.E. and P.N.; writing—original draft preparation, I.E. and P.N.; writing—review and editing, P.N., R.G. and N.B.; visualisation, I.E. and P.N.; supervision, P.N.; project administration, P.N. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data presented in this study are available on request from the corresponding author. The data are not publicly available due to privacy and commercial restrictions. Ethics and privacy: The data were provided by the parking operator's IT contractor as transaction records containing pseudonymous client and session identifiers, vehicle category, and vehicle registration numbers. Registration numbers were the only potentially identifying field and were deleted on receipt, before any analysis. The pseudonymous client identifier was used solely to construct client cohorts and the cumulative client base (CCB) and was not linked to any external personal data. No names, contact details, or payment-instrument data were accessed. Code availability: analysis code is available from the corresponding author on reasonable request. Map data © OpenStreetMap contributors, used under the Open Database Licence.

Acknowledgments: The authors acknowledge the support of the parking service operator for providing the transaction data used in this study.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

CCB	Cumulative Client Base (ever-active definition)
PCA	Principal Component Analysis
SPI	Supply Pressure Index
OSM	OpenStreetMap
POI	Point of Interest

References

1. Ma, C.; Huang, X.; Li, J. A review of research on urban parking prediction. *J. Traffic Transp. Eng. (Engl. Ed.)* **2024**, *11*, 700–720. <https://doi.org/10.1016/j.jtte.2023.11.004>.
2. Shoup, D. *The High Cost of Free Parking*; Routledge: New York, NY, USA, 2021. <https://doi.org/10.4324/9781351179539>.
3. Pierce, G.; Shoup, D. Getting the prices right: An evaluation of pricing parking by demand in San Francisco. *J. Am. Plann. Assoc.* **2013**, *79*, 67–81. <https://doi.org/10.1080/01944363.2013.787307>.
4. Ottosson, D.B.; Chen, C.; Wang, T.; Lin, H. The sensitivity of on-street parking demand in response to price changes: A case study in Seattle, WA. *Transp. Policy* **2013**, *25*, 222–232. <https://doi.org/10.1016/j.tranpol.2012.11.013>.
5. Qian, Z.S.; Rajagopal, R. Optimal occupancy-driven parking pricing under demand uncertainties and traveler heterogeneity: A stochastic control approach. *Transp. Res. Part B Methodol.* **2014**, *67*, 144–165. <https://doi.org/10.1016/j.trb.2014.03.002>.
6. Arnott, R.; Inci, E. An integrated model of downtown parking and traffic congestion. *J. Urban Econ.* **2006**, *60*, 418–442. <https://doi.org/10.1016/j.jue.2006.04.004>.
7. Katambire, V.N.; Musabe, R.; Uwitonze, A.; Mukanyiligira, D. Forecasting the traffic flow by using ARIMA and LSTM models: Case of Muhima Junction. *Forecasting* **2023**, *5*, 616–628. <https://doi.org/10.3390/forecast5040034>.
8. Song, B.; Zhou, H.; Huang, X.; Ma, W. Collaborative parking vacancy prediction for cities with partial sensors missing. *IEEE Trans. Intell. Transp. Syst.* **2024**, *25*, 11439–11451. <https://doi.org/10.1109/TITS.2024.3377734>.
9. Li, J.; Ma, H.; Li, X.; Ma, R. Semantics-Aware Spatial-Temporal Dynamic Graph Transformer Network for On-Street Parking Occupancy Prediction. *IEEE Trans. Intell. Transp. Syst.* **2026**, *27*, 842–856. <https://doi.org/10.1109/TITS.2025.3629329>.
10. Ma, H.; Li, X.; Li, J.; Ma, R.; Pedrycz, W. Improving Park-and-Ride Occupancy Prediction with Feature-Augmented Cross-Variate and Temporal Transformer. *IEEE Trans. Intell. Transp. Syst.* **2026**, *27*, 8835–8848. <https://doi.org/10.1109/TITS.2026.3665597>.
11. Ye, X.; Wang, J.; Wang, T.; Yan, X.; Ye, Q.; Chen, J. Short-term prediction of available parking space based on machine learning approaches. *IEEE Access* **2020**, *8*, 174530–174541. <https://doi.org/10.1109/ACCESS.2020.3025589>.
12. Nie, Y.; Yang, W.; Chen, Z.; Lu, N.; Huang, L.; Huang, H. Public curb parking demand estimation with POI distribution. *IEEE Trans. Intell. Transp. Syst.* **2022**, *23*, 4614–4624. <https://doi.org/10.1109/TITS.2020.3046841>.
13. Vlahogianni, E.I.; Kepaptsoglou, K.; Tsetos, V.; Karlaftis, M.G. A real-time parking prediction system for smart cities. *J. Intell. Transp. Syst.* **2016**, *20*, 192–204. <https://doi.org/10.1080/15472450.2015.1037955>.
14. Yang, S.; Ma, W.; Pi, X.; Qian, S. A deep learning approach to real-time parking occupancy prediction in transportation networks incorporating multiple spatio-temporal data sources. *Transp. Res. Part C Emerg. Technol.* **2019**, *107*, 248–265. <https://doi.org/10.1016/j.trc.2019.08.010>.
15. Gkolias, K.; Vlahogianni, E.I. Convolutional neural networks for on-street parking space detection in urban networks. *IEEE Trans. Intell. Transp. Syst.* **2019**, *20*, 4318–4327. <https://doi.org/10.1109/TITS.2018.2882439>.
16. Geng, Y.; Cassandras, C.G. New ‘smart parking’ system based on resource allocation and reservations. *IEEE Trans. Intell. Transp. Syst.* **2013**, *14*, 1129–1139. <https://doi.org/10.1109/TITS.2013.2252428>.
17. Zheng, Y.; Rajasegarar, S.; Leckie, C. Parking availability prediction for sensor-enabled car parks in smart cities. In *Proceedings of the IEEE Tenth International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP)*, Singapore, Singapore, 7–9 April 2015; pp. 1–6. <https://doi.org/10.1109/ISSNIP.2015.7106902>.
18. Al-Turjman, F.; Malekloo, A. Smart parking in IoT-enabled cities: A survey. *Sustain. Cities Soc.* **2019**, *49*, 101608. <https://doi.org/10.1016/j.scs.2019.101608>.
19. Sun, H.; Huang, X.; Ma, W. Beyond prediction: On-street parking recommendation using heterogeneous graph-based list-wise ranking. *IEEE Trans. Intell. Transp. Syst.* **2024**, *25*, 5892–5903. <https://doi.org/10.1109/TITS.2023.3336808>.

20. Fotheringham, A.S.; O’Kelly, M.E. *Spatial Interaction Models: Formulations and Applications*; Kluwer Academic Publishers: Dordrecht, The Netherlands, 1989.
21. Millard-Ball, A.; Weinberger, R.R.; Hampshire, R.C. Is the curb 80% full or 20% empty? Assessing the impacts of San Francisco’s parking pricing experiment. *Transp. Res. Part A Policy Pract.* **2014**, *63*, 76–92. <https://doi.org/10.1016/j.tra.2014.02.016>.
22. Yalvac, E.; Kay, M.G. Synthetic demand flow generation using the proximity factor. *Forecasting* **2025**, *7*, 14. <https://doi.org/10.3390/forecast7010014>.
23. Geng, Y.; Cassandras, C.G. A new ‘smart parking’ system infrastructure and implementation. *Procedia Soc. Behav. Sci.* **2012**, *54*, 1278–1287. <https://doi.org/10.1016/j.sbspro.2012.09.842>.
24. Belaroussi, R.; Delhoum, Y. Forecasting daily activity plans of a synthetic population in an upcoming district. *Forecasting* **2024**, *6*, 378–403. <https://doi.org/10.3390/forecast6020021>.
25. Tobin, J. Estimation of relationships for limited dependent variables. *Econometrica* **1958**, *26*, 24–36. <https://doi.org/10.2307/1907382>.
26. Talluri, K.T.; van Ryzin, G.J. *The Theory and Practice of Revenue Management*; Springer: New York, NY, USA, 2004. <https://doi.org/10.1007/b139000>.
27. Gammelli, D.; Peled, I.; Rodrigues, F.; Pacino, D.; Kurtaran, H.A.; Pereira, F.C. Estimating latent demand of shared mobility through censored Gaussian processes. *Transp. Res. Part C Emerg. Technol.* **2020**, *120*, 102775. <https://doi.org/10.1016/j.trc.2020.102775>.
28. Hüttel, F.B.; Peled, I.; Rodrigues, F.; Pereira, F.C. Modeling censored mobility demand through censored quantile regression neural networks. *IEEE Trans. Intell. Transp. Syst.* **2022**, *23*, 21753–21765. <https://doi.org/10.1109/TITS.2022.3190194>.
29. Zhang, Z.; Zhang, F.; Liu, W.; Yang, H. On the service differentiation for parking sharing. *Transp. Res. Part C Emerg. Technol.* **2025**, *170*, 104915. <https://doi.org/10.1016/j.trc.2024.104915>.
30. Alampoor, A.; Shafahi, Y.; Levin, M.W. Developing urban parking supply and pricing policies in response to automated vehicle market penetration. *Transp. Res. Part C Emerg. Technol.* **2025**, *181*, 105363. <https://doi.org/10.1016/j.trc.2025.105363>.
31. Kagan, D.Z.; Rylov, A.A. Cluster analysis Russian Federation regions by demand for transport services. *Digital Solutions and Artificial Intelligence Technologies* **2025**, *1*, 77–88. (In Russian) <https://doi.org/10.26794/3033-7097-2025-1-3-77-88>.
32. Yatskin, D.V.; Kochkarov, A.A.; Okuneva, E.A. Assessment of the structural stability of transport and logistics systems based on graph models and the percolation coefficient. *Digital Solutions and Artificial Intelligence Technologies* **2026**, *2*, 52–62. (In Russian) <https://doi.org/10.26794/3030-7097-2026-2-1-52-62>.
33. Chupin, A.; Morkovkin, D.; Bolsunovskaya, M.; Boyko, A.; Leksashov, A. Techno-economic sustainability potential of large-scale systems: Forecasting intermodal freight transportation volumes. *Sustainability* **2024**, *16*, 1265. <https://doi.org/10.3390/su16031265>.
34. Kurganov, V.; Gryaznov, M.; Aduvalin, A.; Polyakova, L.; Dorofeev, A. Analysis of regulation of costs for operating buses in a transport company. *Sustainability* **2024**, *16*, 7274. <https://doi.org/10.3390/su16177274>.
35. Dorofeev, A.; Altukhova, N.; Filippova, N.; Pashkova, T.; Ponomarev, M. Development of transportation management system with the use of ontological and architectural approaches to ensure trucking reliability. *Sustainability* **2020**, *12*, 8504. <https://doi.org/10.3390/su12208504>.
36. Chupin, A.; Sorokin, A.; Veselko, A.; Morkovkin, D.; Pishchik, V.Y.; Alekseev, P.V. Sustainable financing for transport infrastructure: An integral approach for the Russian Federation. *Sustainability* **2024**, *16*, 3108. <https://doi.org/10.3390/su16083108>.
37. Yukhno, A.S. Data-driven governance as an emerging paradigm in public administration: Conceptual foundations. *Management Sciences* **2026**, *16*, 18–28. (In Russian) <https://doi.org/10.26794/2304-022X-2026-16-3-18-28>.
38. Mei, Z.; Murzin, A.D. Intelligent decision support system for enterprise management. *Management Sciences* **2026**, *16*, 94–103. (In Russian) <https://doi.org/10.26794/2304-022X-2026-16-3-94-103>.
39. Boeing, G. OSMnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Comput. Environ. Urban Syst.* **2017**, *65*, 126–139. <https://doi.org/10.1016/j.compenurbsys.2017.05.004>.
40. Saki, S.; Hagen, T. What drives drivers to start cruising for parking? Modeling the start of the search process. *Transp. Res. Part B Methodol.* **2024**, *188*, 103058. <https://doi.org/10.1016/j.trb.2024.103058>.
41. Eryomin, I.R.; Nikitin, P.V. Analysis of paid parking service data to create an effective pricing system (using the example of Vladivostok). *Modeling, Optimization and Information Technology* **2024**, *12*. (In Russian) <https://doi.org/10.26102/2310-6018/2024.45.2.015>.
42. Lautso, K. Mathematical relationships among parking characteristics and revising and reduction methods of parking field survey information. *Transp. Res. Part B Methodol.* **1981**, *15*, 73–83. [https://doi.org/10.1016/0191-2615\(81\)90034-5](https://doi.org/10.1016/0191-2615(81)90034-5).

43. Dijkstra, E.W. A note on two problems in connexion with graphs. *Numer. Math.* **1959**, *1*, 269–271. <https://doi.org/10.1007/BF01386390>.
44. Huff, D.L. Defining and estimating a trading area. *J. Mark.* **1964**, *28*, 34–38. <https://doi.org/10.1177/002224296402800307>.
45. Jolliffe, I.T. *Principal Component Analysis*, 2nd ed.; Springer: New York, NY, USA, 2002. <https://doi.org/10.1007/b98835>.
46. Breiman, L.; Friedman, J.H.; Olshen, R.A.; Stone, C.J. *Classification and Regression Trees*; Wadsworth: Belmont, CA, USA, 1984.
47. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **2011**, *12*, 2825–2830. Available online: <https://jmlr.org/papers/v12/pedregosa11a.html> (accessed on 17 October 2026).
48. Hoerl, A.E.; Kennard, R.W. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* **1970**, *12*, 55–67. <https://doi.org/10.1080/00401706.1970.10488634>.
49. Cameron, A.C.; Trivedi, P.K. *Regression Analysis of Count Data*, 2nd ed.; Cambridge University Press: Cambridge, UK, 2013. <https://doi.org/10.1017/CBO9781139013567>.
50. Cox, D.R. Regression models and life-tables. *J. R. Stat. Soc. Ser. B Methodol.* **1972**, *34*, 187–202. <https://doi.org/10.1111/j.2517-6161.1972.tb00899.x>.
51. Mantouka, E.; Fafoutellis, P.; Vlahogianni, E.I. Deep survival analysis of searching for on-street parking in urban areas. *Transp. Res. Part C Emerg. Technol.* **2021**, *128*, 103173. <https://doi.org/10.1016/j.trc.2021.103173>.
52. Sun, Z.; Wu, J.; Wang, F.; Tian, Z.; He, Q. Survival analysis of parking duration in urban commercial areas: A case study in Zhengzhou, China. *PLoS ONE* **2025**, *20*, e0323207. <https://doi.org/10.1371/journal.pone.0323207>.
53. Administration of Vladivostok. Resolution No. 1458 of 30 May 2025 Amending Resolution No. 2426 of 11 October 2022 on Territorial and Tariff Zones of Paid Public Parking and Parking Fees. 2025. (In Russian) Available online: <https://base.garant.ru/412091000/> (accessed on 17 October 2026).
54. Geary, R.C. The contiguity ratio and statistical mapping. *J. R. Stat. Soc. Ser. D Stat.* **1954**, *5*, 115–141. <https://doi.org/10.2307/2986645>.
55. Moran, P.A.P. Notes on continuous stochastic phenomena. *Biometrika* **1950**, *37*, 17–23. <https://doi.org/10.1093/biomet/37.1-2.17>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.