



# Improving Service Quality in Station-Based Bike-Sharing Systems: A Data-Driven Evaluation of a Crowdsourced Rebalancing Strategy

İsmail Enes Parlak<sup>1\*</sup>

<sup>1</sup> Bursa Technical University, Department of Industrial Engineering, 16345, Bursa, Türkiye

## ARTICLE INFO

Received Date: 1/10/2025  
Accepted Date: 3/11/2025

Cite this paper as:

Parlak, İ., E. (2026). Improving Service Quality in Station-Based Bike-Sharing Systems: A Data-Driven Evaluation of a Crowdsourced Rebalancing Strategy. *Journal of Innovative Science and Engineering*, 10(1), 86-95.

\*Corresponding author: İsmail Enes Parlak  
E-mail:enes.parlak@btu.edu.tr

**Keywords:**

Bike-sharing systems  
Fleet balancing  
Operations research  
User incentives  
Discrete-event simulation  
Sustainable mobility

© Copyright 2026 by  
Bursa Technical University. Available  
online at <http://jise.btu.edu.tr/>



The works published in Journal of Innovative Science and Engineering (JISE) are licensed under a Creative Commons Attribution-NonCommercial 4.0 International License.

## ABSTRACT

Bike-sharing systems (BSS) are a key component of sustainable urban mobility, but their effectiveness is often hindered by the operational challenge of fleet rebalancing, which leads to service failures and high costs. This study proposes and evaluates a dynamic, rule-based user-incentive policy designed to mitigate system imbalances. We developed a data-driven discrete-event simulation model, calibrated with over one million real-world trips from the Oslo city bike system, to serve as a virtual laboratory for testing the policy's performance. The simulation results, averaged over multiple replications, demonstrate the policy's significant effectiveness. Compared to a baseline scenario with a dock unavailability rate of 6.27%, the incentive policy with a 75% user acceptance rate reduced this failure rate to a negligible 0.25%, representing a 96% improvement in service quality for returning users. This work validates that crowdsourcing rebalancing efforts through user incentives, evaluated via high-fidelity simulation, is a powerful and practical approach for enhancing the operational efficiency and reliability of station-based BSS.

## 1. Introduction

Bike-sharing systems (BSS) have emerged globally as a cornerstone of sustainable urban mobility, offering a flexible and environmentally friendly solution for short-distance travel and the "first-and-last-mile" problem. However, the very success and utility of these systems give rise to their most significant operational challenge: the rebalancing

problem (Schuijbroek et al., 2017). This problem originates from the natural asymmetry of user demand, which creates severe imbalances in the availability of bikes and docks across the station network. The more popular a BSS becomes, the more imbalanced it gets, requiring continuous intervention to maintain its functionality (Singla et al., 2015).

From the user's perspective, system imbalance leads to two primary types of service failure: empty stations with no bikes available and full stations with no empty docks for returns (Liang et al, 2024). These events cause significant user frustration, degrade the quality of service, and can ultimately lead to a loss of user trust and a decline in ridership. For the BSS operator, rebalancing is a relentless and costly endeavor, often representing the largest single operational cost driver (Schuijbroek et al., 2017). Traditional rebalancing, which relies on a dedicated fleet of trucks to manually relocate bicycles, is not only expensive-costing systems like Capital Bikeshare an estimated \$200,000 per month-but also often inefficient, failing to fully resolve the imbalance despite the massive expenditure (Wang et al., 2018). Furthermore, the use of fossil-fuel-burning trucks to maintain a "green" transportation system creates a fundamental contradiction with the environmental mission of BSS (Singla et al., 2015).

This points to a pressing need for more intelligent, cost-effective, and sustainable rebalancing strategies. One of the most promising paradigms to emerge is user-centric rebalancing, which aims to "crowdsource" the rebalancing effort by providing incentives to influence user behavior (Singla et al.,

2015). Instead of a logistics problem solved by trucks, rebalancing is reframed as a behavioral and economic challenge solved by designing mechanisms to align user actions with the system's goal of maintaining equilibrium. As the complexity of modeling such user-centric policies exceeds the capabilities of purely analytical methods, computer simulation has become the state-of-the-art methodology for rigorously evaluating their effectiveness in a realistic, dynamic environment (Jian et al., 2016).

This paper addresses the dynamic fleet balancing problem in a station-based BSS by proposing and evaluating a user-incentive policy. Grounded in a comprehensive analysis of over one million real-world trips from the Oslo city bike system (Bysykkkel, 2016), we develop a high-fidelity discrete-event simulation model to serve as a virtual laboratory. We use this model to quantify the performance of our proposed heuristic incentive policy, comparing it against a baseline scenario to measure its effectiveness in reducing service failures and improving the overall quality of service. This work contributes to the growing body of literature on user-centric rebalancing by providing a practical, data-driven evaluation of an incentive strategy designed for dynamic, real-world conditions.

**Table 1:** A Comparative Overview of Rebalancing Methodologies.

Ref. No	Rebalancing Type	Approach	Core Methodology	Key Contribution/Finding
(Raviv et al., 2013)	Static	Operator	Mixed-Integer Linear Programming (MILP)	Developed MILP formulations with a user dissatisfaction objective; solved instances up to 104 stations.
(Dell'amico, 2014)	Static	Operator	MILP, Branch-and-Cut	Formulated the problem as a one-commodity PDP; created a standard set of public benchmark instances.
(Liang et al, 2024)	Dynamic	Operator	MILP, Simulation	Provided a comprehensive modeling framework and used simulation to compare the impact of different modeling assumptions.
(Singla et al., 2015)	Dynamic	User	Online Learning, Crowdsourcing	Designed a complete architecture for a user-incentive system using optimal pricing; deployed in the real world.
(Pan et al., 2019)	Dynamic	User	Deep Reinforcement Learning (DRL)	Developed a hierarchical DRL framework to learn optimal pricing policies for incentives.
(Chiariotti et al., 2020)	Dynamic	Hybrid	Optimization, Simulation	Proposed a framework combining operator-based rebalancing and user incentives, showing synergistic benefits.
(Fugushi ge, 2022)	Dynamic	User	Agent-Based Modeling (ABM)	Used ABM to evaluate incentive policies and quantify their impact on broader sustainability goals.
This Research	Dynamic	User	Data-Driven Discrete-Event Simulation (DES)	Proposes a rule-based incentive policy; quantifies a 96% reduction in dock unavailability via high-fidelity DES.

The academic literature on bike-sharing rebalancing has evolved significantly, moving from centralized, deterministic logistics problems to decentralized, dynamic, and behavior-driven systems. The proposed strategies are broadly categorized into two main paradigms: operator-centric repositioning, where a

The operator-centric paradigm is the most established approach. Early research focused on the Static Bicycle Repositioning Problem (SBRP), which treats rebalancing as an offline logistics challenge, typically solved overnight. These models are often formulated as variants of the Vehicle Routing Problem (VRP) and solved using Mixed-Integer Linear Programming (MILP) to find optimal truck routes. Foundational works introduced objective functions to minimize user dissatisfaction (Raviv et al., 2013), while a standard set of benchmark instances was established to enable subsequent algorithmic comparisons (Dell'amico, 2014). As the field matured, research has extended beyond basic MILP to include more sophisticated metaheuristics to handle larger, real-world instances, such as multi-objective particle swarm optimization (MS-MOPSO) and distributed ant colony optimization frameworks that integrate demand forecasting (Fan et al., 2019). Acknowledging the limitations of static models, the research frontier has shifted towards the Dynamic Bicycle Repositioning Problem (DBRP), which addresses rebalancing in real-time amidst uncertain user demand. This requires more advanced methodologies, including stochastic programming to formally model demand as random variables (Dell'amico et al., 2018) and sophisticated "predict-then-optimize" frameworks that combine machine learning models with robust optimization to improve operational revenue and efficiency (Zhou et al., 2022; Li et al., 2024).

A more recent and transformative paradigm shift has reframed the problem from a logistics challenge to a behavioral one by leveraging user-centric rebalancing. This approach aims to "crowdsource" the rebalancing effort by providing incentives to influence user choice. The core idea is to design a mechanism, often based on dynamic pricing, that aligns users' self-interest with the system's global objective of maintaining balance. A foundational contribution was made by designing and deploying complete architecture for a crowdsourcing-based incentive system (Singla et al., 2015). This learning-based approach was later advanced through a hierarchical deep reinforcement learning framework that learns complex, state-dependent pricing policies (Pan et al., 2019). Subsequent research has significantly advanced this paradigm by

dedicated fleet of vehicles is used to move assets, and user-centric rebalancing, where the system actively engages its users to help maintain equilibrium. Table 1 provides a comparative overview of seminal and recent works across these paradigms, mapping the intellectual development of the field.

incorporating more sophisticated behavioral models, such as using Markov-Modulated Poisson Processes to model station occupancy dynamics (Chiariotti et al., 2020), applying gamification strategies based on behavioral psychology (Peláez-Rodríguez, 2024), and developing two-stage incentive mechanisms that consider user preferences for different types of rewards (Wang & Wang, 2021). In parallel, other studies have focused on evaluating these policies through the lens of sustainability, using agent-based modeling to quantify their impact on broader societal goals like the reduction of vehicle miles traveled (Fukushige, 2022).

Underpinning many of these advanced approaches is the increasing use of machine learning and the indispensable role of simulation. For demand prediction, a critical input for both dynamic operator and user-centric models, the field has moved towards advanced techniques like Graph Convolutional Neural Networks (GCNN) that capture complex spatio-temporal dependencies between stations (Lin et al., 2018) and comprehensive evaluations of various multivariate regression techniques (Peláez-Rodríguez, 2024). As the complexity of these dynamic, stochastic, and behavior-driven systems outstrips the capabilities of purely analytical methods, computer simulation has become the state-of-the-art methodology for evaluation. It provides a "virtual laboratory" to rigorously test and compare the emergent, system-level consequences of different policies in a realistic environment before costly real-world deployment (Liang et al., 2024). This methodological frontier continues to advance, with recent work introducing cutting-edge techniques such as differentiable agent-based modeling for highly efficient, gradient-based optimization of dynamic pricing policies (Mitomi et al., 2025).

This review reveals a clear trajectory toward dynamic, learning-based, and user-centric rebalancing strategies that are validated through high-fidelity simulation. Some recent studies have even explored hybrid frameworks that combine traditional operator-based rebalancing with user incentives, demonstrating that a joint optimization can yield synergistic benefits (Chiariotti et al., 2020). Our research is situated precisely at this frontier. By developing a heuristic user-incentive policy and

evaluating its performance with a data-driven, discrete-event simulation model, our work directly addresses a well-defined and important gap. The novelty of our contribution lies in the practical application and quantification of a specific incentive strategy on a simulation model grounded in a large-scale, real-world dataset, providing a clear and replicable analysis of its potential to improve service quality and operational efficiency.

## 2. Methodology

This study employs a quantitative, data-driven approach based on discrete-event simulation to evaluate the effectiveness of a user-incentive policy for bike-sharing fleet balancing. The methodology is structured into five main phases: (1) data acquisition and preprocessing, (2) exploratory data analysis to build a realistic demand model, (3) development of the discrete-event simulation model, (4) design of the heuristic incentive policy, and (5) the experimental design for the comparative analysis of the policy's performance.

### 2.1. Data Source and Preprocessing

The foundation of this research is a publicly available dataset of historical trip data from the Oslo City Bike (Oslo Bysykkel) system. The dataset was obtained from the official open data portal and encompasses all recorded trips for the entire 2024 calendar year (Bysykkel, 2016).


The raw data, provided as 12 monthly CSV files, was first consolidated into a single dataset. The `started_at`

and `ended_at` columns, provided in a mixed format, were standardized and converted to datetime objects using the ISO8601 format specifier. The dataset included a pre-calculated duration column (in seconds), which was used for trip duration analysis. To eliminate outliers and likely data errors (e.g., false starts or bikes taken for maintenance), trips with a duration of less than 60 seconds or more than 2 hours (7200 seconds) were filtered out. This cleaning process resulted in a final dataset containing 1,139,966 valid trips. Finally, for temporal analysis, new features were engineered from the `started_at` timestamp, including `start_hour`, `day_of_week`, and `month`.


### 2.2. Exploratory Data Analysis (EDA) and Demand Modeling

To ensure our simulation accurately reflects real-world system dynamics, a thorough exploratory data analysis was conducted on the cleaned dataset. The primary goal of the EDA was to extract the key temporal and spatial patterns of user demand. This involved three key analyses:

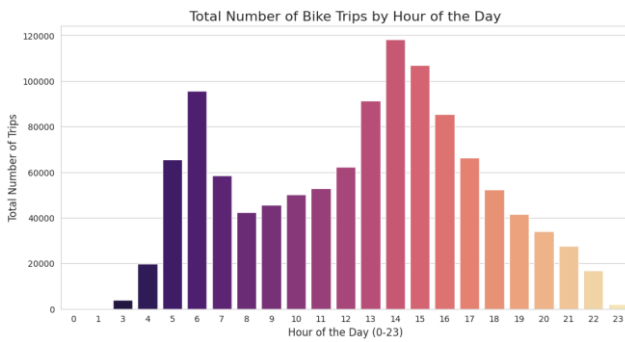
- **Temporal Demand Patterns:** The trip frequency was aggregated by month, day of the week, and hour of the day to quantify seasonal, weekly, and daily patterns. This revealed a strong seasonal dependency with peaks in summer months, a clear distinction between weekday commuter and weekend leisure patterns, and distinct bimodal peaks during morning and evening rush hours (Figures 1-3).



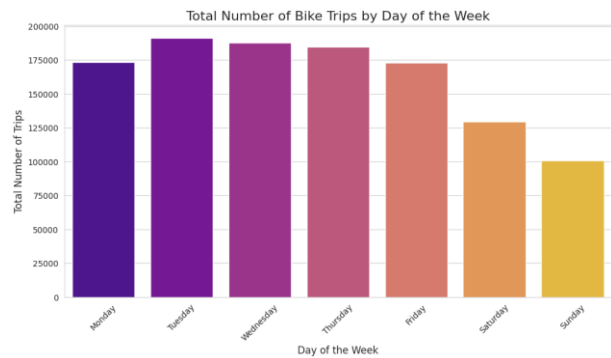
**Figure 1:** Monthly distribution of total bike trips in 2024, illustrating seasonal demand patterns.



**Figure 2:** Distribution of total bike trips by day of the week, highlighting the difference between weekday and weekend usage.



**Figure 3:** Hourly distribution of total bike trips, showing the bimodal peaks corresponding to morning and evening commutes.



**Figure 4:** The 15 most popular origin stations by total number of departing trips in 2024.

- **Spatial Origin Patterns:** The popularity of each station as a trip origin was calculated by counting the total number of trips starting from each unique station. This analysis identified the key demand hotspots within the city (Figure 4).
- **Trip Flow Patterns:** To model trip destinations realistically, an Origin-Destination (OD) matrix was constructed by grouping trips by their start\_station\_id and end\_station\_id. This allowed for the calculation of conditional probabilities, defining the likelihood of a user traveling to a specific destination station given their starting station.

These statistical distributions derived from the EDA served as the foundational logic for the simulation's demand generator.

### 2.3. Discrete-Event Simulation Model

User-centric rebalancing policies involve dynamic and stochastic events, such as uncertain bike arrivals, departures, and trip durations. The complexity of these interactions, combined with the emergent, system-level behavior resulting from individual user decisions, significantly limits the applicability of

purely static or analytical optimization models. For this reason, Discrete-Event Simulation was selected as the foundational methodology. DES allows the system to be modeled as a sequence of events occurring at specific points in time, enabling the model to capture the exact timing and order of critical operational issues, such as a user being unable to return a bike to a full dock. This high-fidelity approach is particularly advantageous as it provides a robust virtual laboratory to test our heuristic incentive policy against the inherent uncertainty (stochasticity) of real-world demand and allows for the precise, quantitative measurement of service failures, which is essential for validating the policy's practical impact.

A discrete-event simulation model of the Oslo BSS was developed using SimPy (v4.1.1), a process-based DES framework in Python. The simulation environment was designed to serve as a virtual laboratory for testing rebalancing policies. The key components of the model are:

- **Environment:** A core SimPy environment managing a virtual clock with a resolution of minutes.

- **Stations:** Each of the unique bike stations from the dataset was modeled as a simply. Container resource with a fixed, uniform capacity of 20 docks. At the start of each simulation run, all stations were initialized with 10 bikes (50% capacity).
- **Trip Generator:** A continuous process that creates user trip events based on the demand models derived from the EDA. At each step in the simulation, the generator decides when to create the next trip (based on a dynamic inter-arrival rate from the hourly demand profile), where the trip should start (based on the station popularity distribution), and where it should go (based on the OD probability matrix).

- **Trip Process:** Each trip is an individual process that executes the following sequence: (1) Arrive at the selected origin station. (2) Request a bike from the station's container. (3) If successful, enter a "travel" state for a duration sampled from the real data's trip duration distribution. (4) Arrive at the destination station and attempt to return the bike to the destination's container.

The conceptual flow of the simulation, illustrating the interaction between the main trip generator loop and the individual trip processes spawned by it, is depicted in Figure 5.

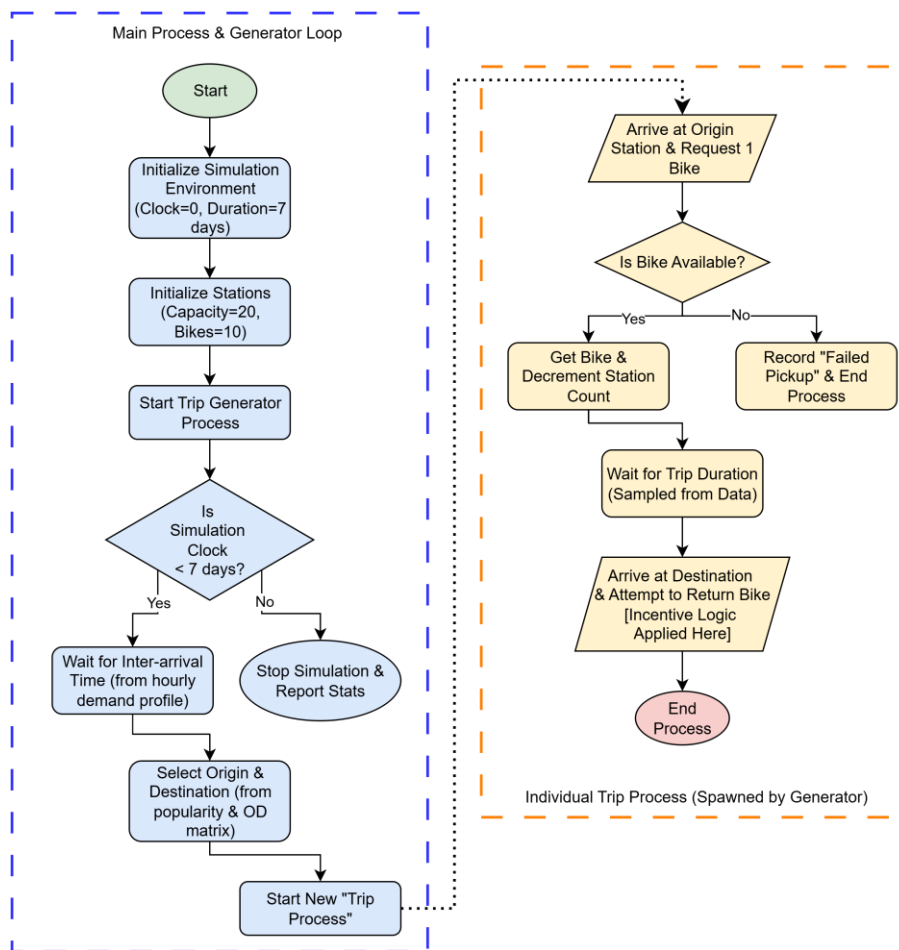


Figure 5: Conceptual flowchart of the discrete-event simulation model.

## 24. Heuristic Incentive Policy

To address the fleet balancing problem, a rule-based heuristic incentive policy was designed and integrated into the simulation. The policy is triggered when a user attempts to return a bike and operates on the following logic:

- **Trigger Condition:** The policy activates if the destination station's bike count has reached a predefined "almost full" threshold (e.g., 90% of capacity).
- **Alternative Search:** The system then searches for a nearby, "almost empty" station (e.g., less than 25%

of capacity). In this study, "nearby" was simplified to a random selection from all qualifying empty stations.

- **User Decision:** If a suitable alternative is found, the user is offered an incentive. The user "accepts" this offer based on a predefined probability ( $P\_ACCEPT\_INCENTIVE$ ), simulating the user's willingness to cooperate.

## 2.5. Experimental Design

To evaluate the policy's effectiveness, a comparative analysis was performed by running the simulation for multiple scenarios over a period of seven virtual days. First, a baseline scenario was established by running the simulation with no incentive policy in effect ( $P\_ACCEPT\_INCENTIVE = 0$ ) to represent the natural, unmanaged state of the system. Subsequently, the incentive policy was activated and tested under varying levels of user cooperation, with the  $P\_ACCEPT\_INCENTIVE$  parameter set to 25%, 50%, and 75%. The performance of each scenario was measured by tracking key performance indicators (KPIs), specifically the total number of failed pickups (due to empty stations) and failed returns (due to full stations), along with their corresponding failure rates.

All simulations were conducted on a Google Colab instance. The computational experiments, including 40 total simulation runs (10 replications for each of the 4 scenarios), were executed on a system equipped with an Intel(R) Xeon(R) CPU @ 2.00GHz (2 vCPUs) and approximately 13 GB of RAM. The total execution time for the entire experimental setup was 4107.57 seconds.

## 3. Results and Discussion

This section presents and discusses the aggregated results from the replicated simulation experiments. Each scenario was run 10 times over a seven-day virtual period, and the averaged results are presented

to ensure statistical robustness and to mitigate the effects of stochastic variability.

The experimental design utilized 10 replications ( $N=10$ ) for each scenario to ensure that the reported average performance metrics were robust against the inherent stochastic variability of the discrete-event simulation model. While a formal statistical power analysis, such as the use of ANOVA or Tukey's HSD test on the results of  $N=10$ , is typically required to prove statistical significance in the strictest sense, the magnitude of the observed differences, particularly the 96% reduction in dock unavailability between the baseline and the 75% acceptance rate scenarios (from 6.27% to 0.25% failure rates), is overwhelmingly large. This dramatic effect size strongly suggests that the performance improvement is a genuine consequence of the incentive policy, making the difference practically and visually significant despite the conservative number of replications. The decision to use  $N=10$  at this stage provides a computationally feasible and highly indicative comparative analysis.

### 3.1. Baseline Scenario: Quantifying System Failure

The baseline scenario, with no incentive policy in effect, was simulated to quantify the natural state of the unmanaged system. The results, averaged over 10 replications, reveal a significant level of service failure. The system consistently produced a bike unavailability rate of 7.15% and a dock unavailability rate of 6.27%. These metrics indicate that in a typical week, a substantial portion of user attempts would fail due to system imbalances, validating the critical need for an effective rebalancing strategy.

### 3.2. Evaluating the Heuristic Incentive Policy

The performance of the heuristic incentive policy was evaluated under three levels of user cooperation (25%, 50%, and 75% acceptance probability). The averaged results, contrasted with the baseline, are presented in Table 2.

**Table 2:** A Comparative Overview of Rebalancing Methodologies.

Acceptance Rate (%)	Avg. Total Trips	Avg. Failed Pickups	Avg. Failed Returns	Pickup Failure Rate (%)	Return Failure Rate (%)	Avg. Incentives Offered	Avg. Incentives Accepted
0	5,813.90	415.80	364.30	7.15	6.27	0	0
25.0	3,965.90	154.20	144.10	3.89	3.63	538.70	133.80
50.0	4,557.00	85.10	60.90	1.87	1.34	508.10	259.30
75.0	5,539.20	104.80	14.10	1.89	0.25	620.10	469.80

The introduction of the user-incentive policy yielded a clear and substantial improvement in system performance across all metrics. A strong dose-response relationship is evident: as user cooperation increases, service failures decrease dramatically.

With a modest 25% acceptance rate, the policy already demonstrates significant impact, reducing the dock unavailability rate by 42% (from 6.27% to 3.63%) and the bike unavailability rate by 46% (from 7.15% to 3.89%). When the user acceptance rate is increased to 50%, the dock unavailability rate is reduced by a remarkable 79% relative to the baseline.

At the 75% acceptance rate, the policy's primary target, dock unavailability, is virtually eliminated, falling from 6.27% to a negligible 0.25%. This represents a 96% reduction in failed returns. Interestingly, while the bike unavailability rate also remains low at this level (1.89%), it does not improve further compared to the 50% scenario. This suggests that while the incentive policy is exceptionally effective at preventing station saturation (full stations), a small residual level of station depletion (empty stations) persists, likely caused by intense, concentrated demand flows during peak hours that even a perfectly balanced return-side cannot fully mitigate. This finding highlights the complex nature of the balancing problem.

The mechanism of this success is clear from the incentive statistics. In the 75% scenario, the system proactively offered an average of 620 incentives per week and successfully diverted nearly 470 users away from full stations, thereby preventing these events from becoming service failures.

### 3.3. Discussion and Implications

The results of this study robustly demonstrate that a data-driven, rule-based user-incentive policy is a highly effective and viable strategy for dynamic fleet balancing. The most significant implication for BSS operators is that even a low level of user cooperation can yield substantial improvements in service quality. This suggests that operators can achieve significant operational gains by implementing even simple incentive schemes that successfully engage a minority of their user base. This crowdsourced approach presents a cost-effective and environmentally sustainable alternative to traditional, capital-intensive rebalancing operations that rely on vehicle fleets.

While the simulation provides strong evidence for the policy's effectiveness, the model includes several simplifying assumptions that represent limitations

and opportunities for future work. These include uniform station capacities, a non-geographical method for selecting alternative stations, and a fixed incentive acceptance probability. Future research could enhance the model's realism by incorporating real-world station capacities, using geographic distance to propose the closest empty station, and developing a dynamic behavioral model where user acceptance depends on the incentive value and required detour.

Building on the highly promising results presented here, future research should focus on rigorously confirming the statistical significance of the policy's benefits. Specifically, we recommend conducting a more extensive simulation campaign, potentially involving  $N \geq 30$  replications per scenario, to allow for formal statistical testing. This expanded data collection would enable the use of techniques like bootstrapping or variance reduction methods to narrow the confidence intervals for the performance metrics, ultimately allowing researchers to confidently assert that the observed differences between the policy and the baseline are statistically significant under all conditions, not just practically evident.

## 4. Conclusion

This paper addressed the critical operational challenge of fleet balancing in station-based bike-sharing systems. By developing a discrete-event simulation model grounded in over one million real-world trips from the Oslo BSS, we quantitatively evaluated the performance of a heuristic user-incentive policy. The results conclusively show that the proposed policy dramatically reduces service failures. Compared to an unmanaged baseline system with high rates of bike and dock unavailability (7.15% and 6.27%, respectively), the incentive policy, with a 75% user acceptance rate, virtually eliminated dock unavailability (0.25%) and significantly reduced bike unavailability (1.89%). This work validates that user-centric strategies, tested via high-fidelity simulation, represent a powerful and practical approach for improving the efficiency, reliability, and sustainability of modern urban mobility systems.

## Article Information

**Financial Disclosure:** The author(s) declare that they have received no financial support for the research, authorship, or publication of this article.

**Authors' Contribution:** The author was responsible for the conception, design, data analysis, model development, and manuscript writing of this study.

**Conflict of Interest/Common Interest:** The author(s) declare no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

**Ethics Committee Approval:** Ethics committee approval was not required for this study as it was based on the analysis of publicly available, anonymized data.

**Declaration of the Author(s):** The author declares that there is no conflict of interest regarding the publishing of this article, that the article has not been published elsewhere, and that it does not include any form of plagiarism. The author has approved the manuscript and agreed with its submission to the *Journal of Innovative Science and Engineering*.

**Acknowledgements:** The authors wish to express their gratitude to Oslo Bysykkel for making their high-quality, historical trip data publicly available. This open data was indispensable for the development and validation of the simulation model presented in this study.

**Data availability statement:** The historical trip data used in this study are openly available in the Oslo Bysykkel Open Data portal, which can be accessed at <https://oslobysykkel.no/en/open-data/historical-website>.

**Artificial Intelligence Statement:** The author bears full responsibility for the content and accuracy of their work, including any use of artificial intelligence (AI) technologies, and confirm that they have read the AI Policy, which is accessible on the journal's website.

---

## References

- Chiariotti, F., Pielli, C., Zanella, A., & Zorzi, M. (2020). A bike-sharing optimization framework combining dynamic rebalancing and user incentives. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, 14(3), 1-30.
- Chiariotti, F., Pielli, C., Zanella, A., & Zorzi, M. (2018). A dynamic approach to rebalancing bike-sharing systems. *Sensors*, 18(2), 512.
- Dell'Amico, M., Hadjicostantinou, E., Iori, M., & Novellani, S. (2014). The bike sharing rebalancing problem: Mathematical formulations and benchmark instances. *Omega*, 45, 7-19.
- Dell'Amico, M., Iori, M., Novellani, S., & Subramanian, A. (2018). The bike sharing rebalancing problem with stochastic demands. *Transportation research part B: methodological*, 118, 362-380.
- Fan, Y., Wang, G., Lu, X., & Wang, G. (2019). Distributed forecasting and ant colony optimization for the bike-sharing rebalancing problem with unserved demands. *PloS one*, 14(12), e0226204.
- Fukushige, T. (2022). Incentive-based Approach to Rebalancing a Dock-less E-Bike-Share System for Sustainability (Doctoral dissertation, University of California, Davis).
- Jian, N., Freund, D., Wiberg, H. M., & Henderson, S. G. (2016, December). Simulation optimization for a large-scale bike-sharing system. In *2016 winter simulation conference (WSC)* (pp. 602-613). IEEE.
- Li, X., Wang, X., & Feng, Z. (2024). Dynamic repositioning in bike-sharing systems with uncertain demand: An improved rolling horizon framework. *Omega*, 126, 103047.
- Liang, J., Jena, S. D., & Lodi, A. (2024). Dynamic rebalancing optimization for bike-sharing systems: A modeling framework and empirical comparison. *European Journal of Operational Research*, 317(3), 875-889.
- Lin, L., He, Z., & Peeta, S. (2018). Predicting station-level hourly demand in a large-scale bike-sharing network: A graph convolutional neural network approach. *Transportation Research Part C: Emerging Technologies*, 97, 258-276.
- Mitomi, T., Makinoshima, F., Makihara, F., & Segawa, E. (2025). Designing Dynamic Pricing for Bike-sharing Systems via Differentiable Agent-based Simulation. *arXiv preprint arXiv:2507.23344*.
- Oslo Bysykkel. (2025). Oslo Bysykkel Open Data: Historical Trip Data [Data set]. Retrieved September 10, 2025, from <https://oslobysykkel.no/en/open-data/historical-website>

- Pan, L., Cai, Q., Fang, Z., Tang, P., & Huang, L. (2019, July). A deep reinforcement learning framework for rebalancing dockless bike sharing systems. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 33, No. 01, pp. 1393-1400).
- Peláez-Rodríguez, C., Pérez-Aracil, J., Fister, D., Torres-López, R., & Salcedo-Sanz, S. (2024). Bike sharing and cable car demand forecasting using machine learning and deep learning multivariate time series approaches. *Expert Systems with Applications*, 238, 122264.
- Raviv, T., Tzur, M., & Forma, I. A. (2013). Static repositioning in a bike-sharing system: models and solution approaches. *EURO Journal on Transportation and Logistics*, 2(3), 187-229.
- Schuijbroek, J., Hampshire, R. C., & Van Hoes, W. J. (2017). Inventory rebalancing and vehicle routing in bike sharing systems. *European Journal of Operational Research*, 257(3), 992-1004.
- Singla, A., Santoni, M., Bartók, G., Mukerji, P., Meenen, M., & Krause, A. (2015). Incentivizing users for balancing bike sharing systems. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 29, No. 1).
- Wang, J., & Wang, Y. (2021). A two-stage incentive mechanism for rebalancing free-floating bike sharing systems: Considering user preference. *Transportation research part F: traffic psychology and behaviour*, 82, 54-69.
- Wang, S., He, T., Zhang, D., Shu, Y., Liu, Y., Gu, Y., ... & Son, S. H. (2018). BRAVO: Improving the rebalancing operation in bike sharing with rebalancing range prediction. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(1), 1-22.
- Zhou, Y., Li, Q., Yue, X., Nie, J., & Guo, Q. (2022). A novel predict-then-optimize method for sustainable bike-sharing management: a data-driven study in China. *Annals of Operations Research*, 1-33.